

# DRAFT

## CMS Paper

*The content of this note is intended for CMS internal use and distribution only*

2013/01/08

Head Id: 119702

Archive Id: 163739:163775M

Archive Date: 2012/05/03

Archive Tag: trunk

## The Detector and Physics Performance of CMS

The CMS Collaboration

### Abstract

The Detector and Physics Performance of CMS

This box is only visible in draft mode. Please make sure the values below make sense.

PDFAuthor: Tejinder Virdee, Giovanni Franzoni

PDFTitle: CMS Paper Template 2006 LaTeX/PdfLaTeX version

PDFSubject: CMS

PDFKeywords: CMS, physics, software, computing

Please also verify that the abstract does not use any user defined symbols



**Contents**

1			
2	1	Introduction . . . . .	2
3	2	Magnet, Infrastructure, Commissioning and Operation . . . . .	2
4	2.1	Infrastructure . . . . .	2
5	2.2	Magnet . . . . .	5
6	2.3	Commissioning . . . . .	9
7	3	Objects Reconstruction and Physics Performance - benchmarked against chal- lenging physics reaction(s) . . . . .	13
8	3.1	Tracking, vertexing and b-tagging . . . . .	13
9	3.2	Electrons/Photons . . . . .	29
10	3.3	Muons . . . . .	29
11	3.4	Jet/MET . . . . .	45
12	3.5	Anomalous Signals in HCAL . . . . .	45
13	3.6	Global event description . . . . .	49
14	4	Level-1 Trigger, HLT . . . . .	49
15	4.1	Level-1 Trigger . . . . .	49
16	4.2	High Level Trigger . . . . .	53
17	5	DAQ . . . . .	61
18	6	Software . . . . .	61
19	6.1	High Level trigger . . . . .	61
20	6.2	Data Handling . . . . .	61
21	6.3	Reconstruction software . . . . .	61
22	6.4	Algorithms improvement for 2012 data taking . . . . .	61
23	6.5	Performance of 2012 reconstruction code . . . . .	62
24	6.6	Alignment and Calibration . . . . .	63
25	6.7	Physics Validation . . . . .	64
26	7	Computing . . . . .	65
27	8	Summary . . . . .	65
28			

## 1 Introduction

T. Virdee

CMS uses a right-handed coordinate system, with the origin at the nominal interaction point, the  $x$ -axis pointing to the centre of the LHC, the  $y$ -axis pointing up (perpendicular to the LHC plane), and the  $z$ -axis along the anticlockwise-beam direction. The polar angle,  $\theta$ , is measured from the positive  $z$ -axis and the azimuthal angle,  $\phi$ , is measured in the  $x$ - $y$  plane.

## 2 Magnet, Infrastructure, Commissioning and Operation

### 2.1 Infrastructure

W. Zeuner

To operate and maintain a complex particle physics detector as CMS, several general services are necessary. The most important ones are discussed in the following.

#### 2.1.1 Detector powering

The CMS detector itself as well as all installations to operate it on site is powered by electrical power. The total steady state consumption of the PT5 site with CMS operating amounts to about 9MW [1]. The total installed power input is about 13.5MW. This power is provided mainly by two networks [2]. The general service power arrives the site through an 18KV line inside the LHC tunnel. On site the 18KV are transformed into 230-400V, 50Hz AC. It is used to operate on-site building services such as lighting, ventilation, elevators and access systems. Depending on the importance for safe operation, the consumers of this network are connected to uninterruptable power supplies (UPS) that allow bridging power outages of 2-30 minutes and that act as filters for transients originating from power glitches and fluctuations in the 400kV system outside the CMS site. The entire system is backed up by a diesel engine 750 kVA electrical power, allowing to continue running services relevant for personnel safety, as emergency lights, ventilation and elevators, for an extended time. The electrical system is operated and monitored by PLCs that require electrical power. For the safe operation of the entire electrical infrastructure a highly redundant network of 48/24V DC is fed from general service power. To ensure continuous operation under almost all circumstances this system is equipped with large battery backed uninterruptable power supplies, allowing running autonomously for several hours in case also the Diesel fails. This system also provides a very limited amount of electricity at 230V AC to operate anti-panic lights in the underground facilities.

The electricity of the second power source, the so-called machine power, arrives via an 18KV surface line from PT6. On site it is transformed into 230-400V AC, available as single phase, three phase and three phase plus neutral return. This provides the necessary flexibility to be able to connect the large variety of powering equipment and electronics needed for a detector of the complexity of CMS. To avoid waveform distortions resulting from the inductive load of the magnets this 18KV system is equipped with static compensators located at PT6. The largest consumers of this power are the cooling stations with 5.5MW and the electronic racks, with up to 2.3MW. The low voltage for the front-end-electronics requires 1MW [1]. Those systems, which are vulnerable to transients from power glitches or to distortions arising from switching on the compensator are connected to battery backed UPS.

### 2.1.2 Detector cooling

The general design philosophy for the CMS infrastructure was that each sub-system has to cool away any excess heat produced by its consumers. The general cooling and ventilation system of the underground caverns should not be additionally charged with heat produced by electrical load. In general this principle was followed by all components. As a consequence, the large consumption of electricity by CMS requires powerful cooling systems. The largest system is the electronic racks dissipating 1.6 MW, followed by the front-end-electronics with 0.8 MW [1]. Underground there are five independent cooling circuits, each of them with their own chiller, pumps and regulation that distribute 18 °C.

### 2.1.3 Gas system

The CMS experiment at the LHC makes use of several gaseous particle detectors for muon identifications as well as for other studies in the high  $\eta$  regions close to the beam pipe. The muon system is employing three different detectors: Cathode Strip Chambers (CSCs), Drift Tubes (DTs) and Resistive Plate Chambers (RPCs). In the high  $\eta$  regions two other detectors are installed TOTEM-T1 and T2, consisting of CSCs and GEMs, respectively. Each type of detector is equipped with a dedicated gas system having the function to provide the suitable gas mixture for the detector operation in terms of mixing ration between the main components and also in terms of impurities, mainly oxygen, nitrogen and water that may be accumulated during extended periods of operation. Each gas system is equipped with a PLC running the control software. The gas systems for all LHC experiments were built according to a common standard to minimize manpower and costs for maintenance and operation. Basically identical modules can be individually configured to satisfy the specific needs of every gaseous particle detector. Most of the gas systems are designed to re-circulate 90% to 95% of the injected gas. As a gas exchange rate between 0.1-0.5 volumes per hour is required for operation and the total detector volumes vary of between 14m<sup>3</sup> and 250m<sup>3</sup> re-circulation is the only way to keep the cost for the primary gas supplies under control. Furthermore this approach reduces to a minimum the consumption of greenhouse gases as, SF<sub>6</sub>, C<sub>2</sub>H<sub>2</sub>F<sub>4</sub> and CO<sub>2</sub>.

The gas systems of CMS are distributed over three different areas at the surface and underground, connected by several hundred meters of pipes. The primary gas supplies, the gas mixing modules, the purifier modules, the exhaust modules and the gas analysis modules are located in a dedicated gas building at the surface. The primary gases used at CMS are Ar, CO<sub>2</sub>, CF<sub>4</sub>, C<sub>2</sub>H<sub>2</sub>F<sub>4</sub>, iC<sub>4</sub>H<sub>10</sub>, SF<sub>6</sub>, Ar/H<sub>2</sub> and N<sub>2</sub>, each provided by two independent supplies that are automatically changed over when one set of gas cylinders is empty. Gases that are liquid at ambient temperature and low pressure (i.e. C<sub>2</sub>H<sub>2</sub>F<sub>4</sub>, iC<sub>4</sub>H<sub>10</sub>) are permanently heated. PLC controlled Mass Flow Controllers are used to mix the primary gases in the correct proportion. After mixing at the surface all gases used at CMS are non-flammable. From the surface building the gases are transported through some hundred meters long pipes to pre-distribution modules located in an underground gas room. From there the gas mixtures are distributed to the different final distributions racks located in the experimental cavern. The underground gas room houses a circulation pump to collect the return gas from the detectors and to send it back to the surface building where it is purified before being re-used. Regulation valves located on the return lines are responsible for the pressure regulation in different regions of the detectors.

Table of key parameters for CMS/Totem gas systems.

### 2.1.4 Compressed dry air and nitrogen supplies

CMS uses nitrogen for inerting, drying to avoid condensation and to steer pneumatic valves. When the detector is closed its inner part flushed with nitrogen. Particularly the inner volumes

of the Pixel detector and the Silicon tracker have to be kept very dry to allow operating at low temperatures; therefore they are particularly carefully sealed. For fire prevention the paraffin end-wall of the Preshower inside the vacuum tank of the magnet has to be in a nitrogen atmosphere. The photo detectors of the ECAL and HCAL are very sensitive to helium, which can be released in case of a fast dump of the solenoid. Therefore they are also permanently flushed with nitrogen. The nitrogen system is backed up by compressed dry air, which kicks in if the nitrogen system fails. As a second backup the inner detector will be flushed by nitrogen from a battery of bottles. For safety reasons the flushing is switched to dry air during maintenance when the yoke is open. Furthermore dry air is used to flush cooling cabinets that are accessible to keep the dew point well below the water condensation limit and to operate pneumatic valves of the cryogenic system of the magnet. The current average consumption is about 105 Nm<sup>3</sup>/h of compressed dry air and about 62 Nm<sup>3</sup>/h of nitrogen. For the future the installation of a large production plant for dry gas is foreseen that can provide up to 400 Nm<sup>3</sup>/h of either dry air or nitrogen, allowing the inner silicon detectors to be operated at temperatures far below 0 °C.

### 2.1.5 Moving system to open and close the detector

The CMS yoke has been constructed in 11 slices perpendicular to the beam axis. The HF towers in front of the yoke can be lowered to the cavern floor and moved into alcoves at the headwalls of the UXC cavern, leaving about 10m to open the yoke on each end of the hall. The central piece carries the solenoid and is bolted to the cavern floor. On each end, three endcap disks and two barrel wheels can be opened and separated, giving relatively quick access to any detector component, including the Pixel detector and the Silicon tracker inside the vacuum tank of the solenoid. The three endcap disks on each end are connected with a movable cable chain allowing to separate any two adjacent disks up to 3.7m. The elements are moved on a system of air pads for any long-range movements. For the final approach below about 10cm and for the locking of the disks together they move on almost friction free grease pads, allowing a precision positioning of the elements to about 1mm. The elements are pulled with a hydraulic strand jack system with 6 jacks. The cavern floor has a slight tilt (1.234%). Therefore, for the uphill move the system has to pull with about 2.5% of the moving mass, which is about 2600t for the three coupled endcap disks. Constant and smooth down hill movements require a pulling as well as a retaining hydraulic. The elements move with the beam pipe in place. The smallest clearance of about 5cm has the YE1 disk when moving over flanges. A sophisticated system of proximity sensors, lasers and tilt sensor allow monitoring the ideal line of movement with a precision of a few millimeters.

### 2.1.6 Operational Experience

In general all infrastructures worked remarkably well from the beginning. The most frequent interruption of operation is caused by power failures. To reduce the impact of power failures and to improve the recovery time the UPS coverage has been continuously improved. This process will continue through the long shutdown, starting end of 2012. First those components were backed up, which suffered damage from sudden cut-offs. They now can ride through power glitches and in case of a real outage they will be ramped down automatically in a controlled way, minimizing possible damage. In its final state the system will allow CMS to ride through outage of a few seconds and shutdown the entire system in a controlled way for outages of longer duration. The highest priority during all power cuts is to keep the magnet running. A fast discharge will heat up the coil to about 80K, requiring 3-4 days cooling before it can be switched back on. Even a slow discharge that cannot be intercepted will lead to at least 8 hours interruption of the magnet operation. The coldbox can continue running for about 30 minutes and afterwards a local Dewar can provide He for couple of hours. The main concern

is to keep the cooling water temperature for the power converters below 27 °C. Major power outages often also affect the cooling systems, either locally or CERN wide. If the local pumps can be restarted quickly the large heat capacity of the circuits slows down the temperature rise, even without cooling on the primary circuit, which is shared by all underground cooling circuits. By quickly switching off all major power consumers of the detector in the large majority of cases a ramp down of the magnet could be avoided. As power failures usually also affect the LHC, they did not yet cause any significant data loss.

In August 2011 a series of, one by one harmless, equipment failures together with some hidden faults in parameter settings and steering software lead to a data loss of about 100pb-1. In short a leak in a safety valve triggered frequent refills of the leaky circuit, causing pressure variations in the filling circuit, which due to a software bug created an alarm stopping the chiller of filling circuit.

Up to now this was the only data loss of CMS caused by a failure of infrastructure. The underground cooling circuits of CMS using demineralized water have a common buffer tank, including the circuit for the magnet infrastructure. Practical experience showed that maintenance and repair work at any of these circuits can disturb the magnet cooling through the connection via the buffer tank. To make the system more robust buffer tank of the magnet circuit will be separated from the other circuits during the coming shutdown. All gas systems are very reliable and stable, sometimes even able to compensate for instabilities due to gas leaks of the detectors. The gas analysis is the most complicated part of the system and it is a key issue for stable long-term operation of the CMS. Maintaining and improving the gas analysis systems has a high priority for the coming years.

The possibility to relatively quickly open the detector to give access to any component with a few weeks is a unique feature of CMS. The opening system has been proven to work reliably. However, moving objects of several thousand tons along the beam pipe with only a few centimeter clearances has a certain risk. The current system requires about 10 highly specialized persons to safely open CMS. The running schema of LHC with several years of operation with only minimum interruption for maintenance makes it particularly difficult to maintain a highly skilled opening team. Therefore attempts are underway to simplify the operation. The movements will become more reproducible, position corrections in particular those perpendicular to the beam axis will become simpler and parts of the monitoring will be automatized. With that the CMS collaboration will maintain the possibility to safely and quickly open the detector whenever necessary.

### 2.1.7 Acknowledgments

Andrea Gaddi, Roberto Guida and Sergei Lusin are gratefully acknowledged for valuable discussions and providing technical details.

## 2.2 Magnet

### B. Cure

The superconducting magnet for the CMS experiment is among the biggest and most powerful magnets in operation to date in high-energy physics. The CMS magnet has a free bore of 6-m diameter and a 12.5 m length, with a maximum stored energy of 2.6 GJ at the design field of 4 T at the interaction point (IP) of the detector. The magnet is a superconducting solenoid with a 10'000 ton iron yoke comprising the barrel wheels and the endcaps. The coil is operated at 4.5 K. The operation field is 3.8 T at the IP, enough to provide the requested resolution for the



physics at LHC, and it offers a gain of about 10 % on the temperature margin, with a margin of 2.1 K at 3.8 T and 1.9 K at 4 T.

### 2.2.1 Operational experience and data

**Field stability.** The magnetic field is precisely measured with high field NMR probes covering a range from 3 T to 4.6 T. The probes are installed within the inner detector volume at 4 locations: two probes on the mid-plane  $Z=0$  on the external radius of the HB detector, plus two probes on the tracker bulkhead. Over the past 6 years, the NMR measurements have been very stable with a standard deviation  $\sigma < 3.10^{-5}$  T. The magnet current ripple is about 1 ppm (18164 0.02 A). No noticeable displacement of the coil has been recorded.

**Magnet cooling time.** The coil cooling time from 300 K to 4.5 K is about 24 days. Following a fast discharge (FD), the coil temperature reaches 67 K in average, and it takes 5 days to cool down the magnet and refill the liquid helium levels in the cryostats. During the year-end shutdown when the annual maintenance takes place on the cryogenics and the water-cooling system, the magnet temperature is left floating and reaches about 90 K in 40 days. Then it takes 7 days to cool down the magnet to 4.5 K (figure 1).




Figure 1: The profiles of the coil maximum and minimum temperatures during the cooling from 90 K to 4.5 K.

**Magnet ramps and electrical characteristics.** The magnet current ramps and discharges on the dump resistor are displayed in 2. The CMS magnet takes advantage of a 20 kA +26V/-23V two-quadrant dc converter allowing both the ramp up and down of the current in between  $\pm 1.5$  A/s, with the possibility to stop the ramp at intermediate current if the operation conditions request so, limiting both the number of magnetic cycles of the magnet and the number of openings of the switch breakers. The coil insulation to ground is excellent, and the last measurement, done in March 2012 on the coil at 4.5 K with both the busbars and the dump resistance connected, has indicated a resistance to ground above 34 M $\Omega$ . The reliability and sturdiness of the magnet powering and discharge circuits have been confirmed over the past years of operation. This is of primary importance to ensure the general detector safety with respect to the very high energy stored in the magnet and its high energy-over-mass ratio, respectively 2.3 GJ and 10.6 kJ/kg when operating the magnet at 3.8 T.

**Magnetic fringe field.** The fringe field around the magnet yoke is not negligible and it affects all the underground areas. Therefore the work in the experimental area is organized according






Figure 2: The magnet current ramp and discharge profiles.

to the status of the magnet. The stray field prevents mostly all the mechanical activities. Figure 3 shows the computed stray field in UXC5.




Figure 3: The CMS magnet stray field in UXC55, with 3.8 T at IP, in the horizontal mid-plane  $Y=0$ , scale in mT, from 0 to 120 mT, with 5 mT increment (courtesy Vyacheslav Klyukhin).

**Magnet availability at 3.8T.** The technical interventions on the magnet have been carefully organized by the CMS Technical Coordination (TC) and the magnet team to prevent any risk of interfering with the data taking. In Table 1, the availability of the magnet is given. The availability of the magnet at 3.8 T is given as a percentage of the total time when the magnet was requested at 3.8 T. The overall magnet working time at 3.8 T is also given, together with the durations of the planned and unexpected stops, as a percentage of the total time from the magnet start at the beginning of the year till its stop before the year-end shutdown.

**Quantity of magnetic and thermal cycles.** The number of magnetic cycles from zero to the nominal magnetic field of 3.8 T is limited as much as possible to preserve the magnet lifetime. As a consequence, the magnet is not systematically stopped when the access to the experimen-

In % time	Year 2010	Year 2011
Beam time loss	none	none
Availability at 3.8T for physics	99.97 %	100 %
Overall working time at 3.8T	99.97 %	84.79 %
Planned magnet stops	11.71 %	12.21 %
Unexpected magnet stops	1.28 %	3.00 %

Table 1: CMS magnet availability in operation during the 2010 and 2011 physics runs.

	2006	2008	2009	2010	2011	2012	Total
Quantity	12	12	10	14	7	2	57
On request	9	9	5	7	4	1	35
Unexpected	3	3	5	7	3	1	22
FD at nominal field	5	1	1	0	2	0	9

Table 2: Summary of the magnetic cycles of the CMS magnet

tal cavern is granted for a short access or a technical stop. The total quantity of the magnetic cycles is summarized in 2.




Figure 4: The CMS magnet thermal cycle number from 4.5 K since the year 2006.

The largest number of magnetic cycles in 2010 is due to the 5 unexpected field reductions performed to solve the coldbox turbine filter regeneration issue. Concerning the thermal cycles, the magnet has been warmed up to room temperature twice, in 2006 and 2008. There have been a total of 9 fast discharges at nominal current, including the surface test commissioning, causing an average temperature rise to 67 K. The magnet temperature also increases when the coldbox stops, either accidentally or for regeneration, causing a temperature rise kept below 60 K. The summary is given in figure 4.

**Effects of the magnetic and thermal cycles of the CMS magnet.** During the magnetic and thermal cycles of the magnet, the stresses and strains appearing in the superconductor high purity aluminum (HPA) stabilizer induce a degradation of its electrical conductivity. A strain of 0.15% is a typical value for the CMS magnet with a central field of 4 T. The effect of cyclic strains from zero to 0.15% on the HPA electrical resistivity was measured on samples. From the data in [3], the Residual Resistivity Ratio (RRR) of the HPA stabilizer is computed and the results are shown in figure 5 together with the RRR measured on the CMS coil. So far only few measurements could be made, as the magnet must be at an average temperature of about 15

K to measure the RRR. The RRR value after 54 cycles is about 1900. The CMS magnet design studies [4] indicated that a RRR of 500 is still good enough for the coil stability.




Figure 5: Computed and measured RRR of the CMS conductor HPA stabilizer as a function of the cycle number at 0.15% cyclic strain.

### 2.2.2 Problems encountered and causing magnetic cycles

So far the coil never quenched untimely. The magnet safety system worked very well and it systematically triggered when the operating conditions were outside the nominal range. All the quench detectors have been working correctly to date.

**CMS cryogenic plant.** Most of the problems met with the cryoplant arise from pollution due to residual traces of either air or humidity. Such a pollution can cause the blockage of the first heat exchanger of the cold box resulting in a loss of refrigeration power. A full regeneration of the cold box at room temperature is then needed with the magnet off and not refrigerated. This operation must be repeated after 5 months of operation in average, and it takes about one week. It is scheduled each year during the period of the LHC technical stops. The pollution causes also a pressure drop that builds up on the cold box turbine filters. The regeneration of these filters can now be made with the magnet at nominal field. It takes only a couple of hour. It was tested in 2008, and repeated recently in 2012. Other issues were due to faulty sensors or components, causing a magnet ramp down. One turbine of the cold box broke in 2009. Complementary spare parts were purchased. It didn't repeat.

**Power cuts.** Only three CERN-wide power cuts caused a stop of the magnet over the past years: one in 2009, one in 2010 and one in 2011.

**Water Cooling.** The water cooling system never directly caused a magnet stop. It mostly affects the helium compressors and the cold box when the services are switch off by power cuts, but it can be at best recovered quickly in typically a couple of hours without field reduction. Only twice it caused a reduction of the magnetic field.

**Powering system.** One slow discharge was due to a faulty sensor, plus three unintentional switch off by the power converter operators.

## 2.3 Commissioning

T. Camporesi

### 2.3.1 Introduction

The global commissioning of a project like CMS can be separated into three main phases:

1. The availability of an acquisition framework for test beam or laboratory tests, developed centrally, which starts integrating the control and data flow concepts of the XDAQ 5. Such 'standalone' acquisition systems became available around 2004. The move from home developed acquisition systems to the CMS acquisition framework was the first step which allowed the members of the various projects to become acquainted with the control and Hardware environment of CMS.
2. The CMS slice test performed in 2006 at the occasion of the first magnet test, performed on the surface to qualify the magnet ([5]) was the first occasion where most sub detector and the early instantiation of the trigger system came together. This test down under severe time constraints and with just a few elements of the sub detectors was an important step towards forming the CMS team which would work for the subsequent years to commission the detector. In many areas shortcuts and patches had to be implemented to allow collecting the first set of CMS cosmic ray data, but the test allowed the identification of the key areas in the DAQ and trigger frameworks in which needed work and improved functionality. The first rudimentary Slow control system and safety interlocks were also deployed and tested at this time.
3. The lowering of the CMS detector elements into the experimental cavern and their connection to the final readout electronics at LHC point 5 started the final phase of the commissioning of CMS: the first element available to be readout was the one of the forward Calorimeters (HF-) early in 2007.

In this paper we will report only about the third phase. The way CMS was installed makes the commissioning phase completely entangled with the installation phase: commissioning progressed in conjunction, and sometime in competition, with installation. The last element of CMS to be installed was the second ECAL end cap calorimeter in 2009 and its integration in the CMS readout and trigger might be considered to mark the end of the commissioning of CMS and the start of the Running of CMS. The activities during this period can be separated into three main categories which had at times different and possibly contrasting needs with respect to service and infrastructure availability. In the following we use the terms "Global" or 'Central' to identify activities which involved more than 1 subsystem and 'Local' or 'Stand-alone' to identify commissioning activities of a single subsystem.

1. The commissioning of sub detector elements in 'stand-alone' mode: most sub detector teams worked shifts from the moment their detector elements were connected to the readout. Each subsystem had brought their software and control environment from the laboratory or test beam experience. The standalone readout of the detectors was easily brought to the CMS environment. During this phase each CMS partition had its own Local Trigger Controller board which allowed development of control sequences for local data taking similar to the ones needed for central data taking. Clearly local runs were incompatible with central runs and were requiring dedicated services like power and cooling and had safety constraints.
2. The central DAQ setup and Global trigger system had integrated ways to exercise and stress-test their setups in stand-alone mode: e.g. the DAQ could read fake data frames from the FRLs and the trigger could inject patterns into the global trigger sub-units. Such

activities were incompatible with global data taking and required cooling and power in the Service cavern and control rooms.

3. Global data taking was scheduled regularly: during this period one or more sub detector (or parts of) were being triggered and readout centrally: during this period the constraints on the availability of services, safety concerns were maximal.

The real challenge for the whole commissioning period 2007-2009 has been trying to operate detectors or part of them while the whole CMS detector, infrastructure and services were being installed: most of the time there were more than 60 people working in the experimental and service cavern doing mechanical and electrical work. A given element of the detector could be commissioned once connected to the readout electronics. In order to power the detector, front end electronics and readout electronics one needed to have all services connected, detector protection system working, no co-activity near any of the powered elements in order to prevent danger for personnel and/or equipment. Detailed planning was available, but maximum flexibility was needed on the commissioning side as plans had to adapt to all sort of issues arising from the possible problems arising on the technical coordination side. Daily technical coordination meetings took place in order to coordinate activities: most of time installation schedule would take priority, but it was soon realized that some aspect of the central commissioning effort needed to be guaranteed and effort was put in place to schedule activities around some key commissioning events. For example any effort was made to allow scheduled the Global runs, described below: such events required the presence of a lot of experts from each subsystem, some of which flying in from overseas and were key to the progress of global commissioning.

### 2.3.2 Commissioning Strategy

The CMS commissioning goal was defined to be ready when LHC would start delivering data to collect data efficiently, and of good quality and ship it to Tier0 for processing. The metric to evaluate our progress was based on efficiency for data taking, quality of the data collected and effective data flow. As the deployment of the DAQ and trigger infrastructure, software and hardware, the installation of the CMS detectors and services and the local certification of parts of the detectors was taking place concurrently the strategy of the commissioning coordination was based on

1. First :commission functionality, possibly a very limited part of the detector.
2. Second: increase complexity. More and more of given sub detector, commission operation of more than one sub detector, increase trigger rate, increase trigger complexity and so on.
3. Third: increase data taking reliability. While in the first attempts we could tolerate lifetimes below 70% the final goal was to be alive for more than 90% of the time.
4. Fourth: increase data quality. For example, while at the beginning we could tolerate having part of the data from parts of sub detectors with data corruption, the goal was to be close to 100% physics quality data for when the LHC would start.

The way we implemented in practice this strategy was through

1. Scheduling regular Global runs, widely advertised as milestones in the collaboration. This global runs would take place  $\approx$  at the end of every month (would be dubbed Global



Run at End of X, GREX, where X stand for the initial of the month) and would last typically 2 or 3 days. The goals from each Global run would be set in the week ahead of the run, aiming to improve in any of the four commissioning strategy categories, compared to the previous global run. These Global runs were 'compulsory' for any system having any part of the detector ready to take data.

2. From the moment most of the sub detectors became available for readout (late summer 2008) we started setting aside a daily slot every 2 weeks called mid-week global runs. Such runs were meant on one hand to prepare the monthly global runs and on the other were provided as a service to sub-system who wanted to debug/commission changes to their hardware or Software setup.

In these runs we would exercise the whole detector control and data flow chain: each run would go through standard Slow Controls and Fast configuration sequences, would have the whole trigger (Level 1 and HLT) enables, possibly with limited complexity. Since the first run we had online process monitoring and real time Data Quality Monitoring and data would be shipped through the standard data link to the tier 0 for processing and we would monitor the quality of reconstructed data. Of course in all of these we started with very limited functionality/complexity and we increased it at every global run. For example the very first globular un lasted 8 hours and was reading partially one of the HF calorimeters, triggering on random triggers and monitoring the level of the pedestals and verifying that the reconstructed data at Tier0 matched the online monitored one. Already from the second global run we we also exercising shipping data to the rest of the GRID. The pace at which we increased complexity can be seen in Figure 6 In the week long Global run on November 2007 we collected 10 million




Figure 6: The increase of complexity of the Global Runs in 2007 and 2008: the bars are color coded to represent the fraction of a given sub detector which was involved in the global run

triggers for a total of 3 Terabytes of accumulated data (raw + reconstructed), in the long run in November 2008 we would collect 10 Million events in 8 hours and accumulate 2 terabytes of data per hour. During 2008 and 2009 in several occasions we planned some major commissioning event, always timed with respect to the LHC operation, namely the start of circulating beams in the accelerator, signaling the end of the LHC commissioning periods and anticipated collision data. These were extended runs, with 24/7 data taking planned to last several weeks (in one case 2 months): we named them depending on whether the Magnet was off ( Cosmic-ray RUn at Zero Tesla, CRUZET) or on (Cosmic-ray Run at Four Tesla, CRAFT). Noticeable achievements during the CRUZET runs of the fall 2008:

1. Verified the trigger rate capability of our system ( designed for 100 KHz Level 1 and 100 Hz logging rate out of the HLT): by injecting random triggers on top of few hundred Hz cosmic ray triggers we could reliably operate CMS at 80 KHz logging up to 300 Hz of data.

2. Commissioned special calibration sequences, like interleaving 100 Hz of laser shots to monitor the ECAL transparency to regular data taking (such special trigger would be operated during the orbit gap of the emulated LHC orbit cycle)
  3. rehearsed procedures for the remote prompt analysis teams from each sub detector
  4. Monitored continuously the noise behavior of sub detectors: lots was learned about the noise behavior of the various detectors: for some of them the CMS cavern environment had to be understood interns of founding. For the calorimeters, the availability of the central calorimeter trigger allowed for the first time self triggering and identification of certain noise sources.
- These extended runs ended up exceeding our expectations: we collected each year several billions of cosmic ray tracks with the full installed detectors operational. These data sample allowed us to start data taking at LHC with a well aligned and calibrated detector, thus allowing very quick understanding of the detector response and quick production of physics results.
5. The systematic shipping of the data to the CERN Analysis Facility (CAF) and the Tier0 for prompt reconstruction allowed debugging of the Reconstruction and prompt analysis workflows

### 2.3.3 The CRAFT run: allowing CMS ready to take effectively and publish data quickly

## 3 Objects Reconstruction and Physics Performance - benchmarked against challenging physics reaction(s)

### 3.1 Tracking, vertexing and b-tagging

– I. Tomalin

At the LHC design luminosity of  $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ , about 1000 particles from more than 20 overlapping proton-proton interactions will traverse the CMS tracker, at each LHC bunch crossing. Track finding in such a high occupancy environment is immensely challenging. It is not easy to obtain a high tracking efficiency, whilst at the same time keeping the fraction of fake tracks low. In addition, the tracking code must run sufficiently fast that it can be used not only for offline event reconstruction (of  $\approx 10^9$  events per year), but also for the CMS High Level Trigger.

The physics goals of CMS [6, 7] place strong requirements on the performance of the tracking. Searches for high mass dilepton resonances demand that it should have good momentum resolution for transverse momenta  $p_T$  of up to 1 TeV/c. At the same time, efficient reconstruction of very soft tracks with,  $p_T \ll 1 \text{ GeV}/c$ , is needed for studies of hadron production rates and to obtain optimum jet energy resolution with particle flow techniques [8]. In addition, it must be possible to resolve very close tracks, such as those from 3-prong tau decay. And excellent impact parameter resolution is needed for a precise measurement of the primary and secondary vertex positions and for  $b$ -jet identification.

This section describes the algorithms used for track finding, primary and secondary vertex reconstruction, and  $b$  tagging. The performance achieved by these algorithms is shown. In addition, since none of this would be possible without a precise measurement of the alignment of the Tracker modules, we begin by describing how that is done.



### 3.1.1 The CMS Tracker

The CMS silicon Tracker [9] occupies a cylindrical volume of length 5.8 m and diameter of 2.5 m, whose axis is approximately directed along the LHC beam line. It is immersed in a co-axial magnetic field of 3.8 T provided by the CMS solenoid. A schematic drawing of the CMS tracker is shown in Fig. 7. It has an acceptance extending up to a pseudo-rapidity  $|\eta| < 2.5$ .




Figure 7: Schematic cross section through one-quarter of the CMS tracker. Green lines show layers of the Pixel Tracker and blue (red) lines show Strip Tracker layers with (without) the back-to-back modules which deliver stereo hits.

The Pixel Tracker consists of 3 barrel layers at radii from 4.4–10.2 cm, and two pairs of endcap disks at  $z = \pm 34.5$  and  $\pm 46.5$  cm. In total, its 1440 modules cover an area of about  $1 \text{ m}^2$  and have 66 million pixels. They provide 3-D measurements of the positions of the ‘hits’ arising from the interaction of charged particles with its sensors. The Strip Tracker has 15 148 silicon modules, which cover an area of about  $198 \text{ m}^2$  and have 9.3 million strips. It is divided into several sub-systems. The barrel is made of the Inner Barrel (TIB) and small radii and the Outer Barrel (TOB) at large radii. Taken together, these have 10 layers in total. Each endcap is made up of the small Inner Discs (TID) and the larger Endcap (TEC) disks, which together comprise 12 layers.

The modules in the first two layers of the TIB and of the TOB, as well as rings 1 and 2 of the TID and 1, 2 and 5 of the TEC, carry a second strip detector module, which is mounted back-to-back to the first, with a stereo angle of 100 mrad, in order to provide a measurement of the second coordinate ( $z$  in the barrel and  $r$  on the disks), and so allow 3-D hit position measurement.

The Tracker has a relatively high material budget. Depending on the pseudo-rapidity  $\eta$ , this varies from 1.4–1.9 (0.14–0.55) radiation (hadronic interaction) lengths. It is smallest for  $|\eta| < 1$  and largest for  $|\eta| \approx 1.4$ .

### 3.1.2 The alignment of CMS Silicon Tracker

The alignment of the Tracker modules with the desired precision is a necessary step in the commissioning of the detector and plays a crucial role in reaching the physics goals of CMS. Design specifications indicate that the tracking must reach a resolution on the transverse momentum

( $p_T$ ) of 1.5% (10%) for muons of momentum of 100 (1000) GeV/c [6]. Therefore, in order to fully exploit the single hit resolution, which can be as good as 10  $\mu\text{m}$  in the pixel sensors and 14  $\mu\text{m}$  in the strip sensors, the positions of the sensors must be known to a precision of a few micrometers. This can best be achieved by track-based alignment algorithms.

Although the result obtained using cosmic ray tracks only as alignment input [9] has been excellent and was instrumental for the early physics program of CMS, the alignment was still not achieving the final level of accuracy and control of systematic distortions desired, due to limited statistics and the inability of cosmic rays to fully constrain all degrees of freedom. With the inclusion of the large statistics of tracks from pp collisions provided by the LHC, the goals of the alignment of the CMS tracker have been to reach the ultimate statistical resolution in all regions of the tracker and an extensive control of the relevant systematic distortions biasing reconstructed track parameters and thus affecting physics performances.

### 3.1.2.1 The 2011 Tracker alignment strategy

Track-hit residual distributions are generally broadened if the assumed positions and orientations of the silicon modules differ from the true ones. Following the least squares approach, alignment algorithms minimise the squares of normalised residuals, summing over many tracks. If the (hit or virtual) position measurements  $m_{ij}$  with uncertainties  $\sigma_{ij}$  are independent, the minimised function is

$$\chi^2(\mathbf{p}, \mathbf{q}) = \sum_j^{\text{tracks}} \sum_i^{\text{measurements}} \left( \frac{m_{ij} - f_{ij}(\mathbf{p}, \mathbf{q}_j)}{\sigma_{ij}} \right)^2 \quad (1)$$

where  $f_{ij}$  is the track model prediction at the position of the measurement, depending on the alignment ( $\mathbf{p}$ ) and track ( $\mathbf{q}_j$ ) parameters. In a global fit approach as implemented in the MILLEPEDE II program [10, 11]  $\chi^2(\mathbf{p}, \mathbf{q})$  is minimised after linearising  $f_{ij}$  and the alignment parameters are determined.

The improved track model [12, 13] allows also for a better treatment of multiple scattering effects, achieved by increasing the number of parameters for a charged particle in the magnetic field to  $n_{\text{par}} = 5 + 2n_{\text{scat}}$ , e.g. adding two deflection angles for each of the  $n_{\text{scat}}$  thin scatterers.

The alignment of the CMS tracker in 2011 was performed exploiting this global fit approach with the improved track model, using as input data collected until end of June, about 15 million loosely selected isolated muon tracks, 3 million low momentum tracks, 3.6 million cosmic ray tracks (collected between LHC fills, during collisions and before collision data taking) and 375 thousand muon track pairs from Z boson decays.

### 3.1.2.2 Alignment precision

The estimation of the achieved statistical precision of the 2011 alignment was estimated via track-based validation. The input to the validation are isolated muon tracks with a transverse momentum of  $p_T > 40$  GeV and at least ten hits in the Tracker. The tracks are refitted taking the new determined module positions into account. Hit residuals are determined with respect to the track prediction, which is obtained without using the hit in question to avoid any correlation between hit and track. From the residual distribution of the unbiased hit residuals in each module, the median is taken and histogrammed for all modules in a detector subsystem. The median is relatively robust against stochastic effects from multiple scattering, and thus the distribution of medians of residuals (DMR) is taken as a appropriate measure of alignment accuracy. Only modules comprising at least 30 entries in their residual distribution are consid-

ered.

**Track residuals in the Pixel detector.** Compared to the alignment with cosmic rays alone [9], the most striking improvements are observed in the end caps of the pixel tracker, where the addition of tracks from collision events leads to a huge boost of statistics, especially for innermost parts of the tracker. An example of corresponding DMR is shown in the left plots of Fig. 8, separately for the  $u$  and  $v$  coordinates; their RMS is well below  $3 \mu\text{m}$  in both directions, compared to about  $13 \mu\text{m}$  in the cosmics-only alignment. These numbers are only slightly larger than the ones obtained in simulation without any misalignment, which are between  $1\text{--}3 \mu\text{m}$ , and far below the expected hit resolution.




Figure 8: Left: Distributions of the medians of the residuals, for the pixel tracker end cap modules in  $u$  coordinate. Shown in each case are the distributions after alignment with 2011 data, in comparison with simulations without any misalignment and simulation tuned to reproduce the misalignment after the 2011 alignment procedure. Right: Day-by-day value of the relative longitudinal shift between the two half-shells of BPIX as measured with the primary vertex residuals.

**Time dependent corrections to the Pixel high structures** In addition, unbiased track-vertex residuals are also used to monitor the position of the two pixel half barrels relative to each other. Each primary vertex is refitted after removal of one of its tracks. This is repeated for each track of the vertex. The track-vertex residuals  $\Delta z$  along the beam line are averaged as a function of the polar angle  $\phi$  of the track. A difference of the mean values for tracks stemming from the one half barrel or the other indicates a relative misplacement. Jumps of up to  $30 \mu\text{m}$  are seen before alignment.

After the alignment with time dependent parameters for the positions and orientations of large pixel structures, the remaining half barrel separations are well below  $10 \mu\text{m}$ , a value that has no effect on the alignment sensitive  $b$ -tagging algorithms.

### 3.1.2.3 Sensor and module shape parametrization

In the CMS software, the module translations  $u, v, w$  as well as the rotations  $\alpha, \beta, \gamma$  around these axes are defined in the local reference system of the module and determined by the alignment procedure. This assumes flat sensor surfaces. In reality, however, the surfaces of the sensors are not flat. To take this effect into account, for each sensor the sum of second order modified Legendre polynomials has been introduced [14] in order to parametrise the sensor surface. The curvatures of the sensor surfaces are referred to as *bows*. Furthermore, all TOB modules and the

TEC modules at radii  $r > 60$  cm consist of two individual daisy-chained sensors. With respect to alignments of the CMS tracker performed in earlier years, the treatment of these double sensor modules has been improved by allowing the separate determination of the alignment parameters for both sensors. This improvement is referred to as the determination of *kinks*. The parametrization with polynomials describes the sensors very well. This can be seen in Fig. 9 (left), where the offsets  $\Delta w$  are calculated from the residuals in  $u$  and the track angle  $\psi$  from the sensor normal in the  $uw$  plane. For an alignment with the flat sensor assumption, a parabolic shape is seen that vanishes taking into account the additional parameters. High mo-




Figure 9: Left: Distributions of the weighted means of the  $\Delta w = \Delta u / \tan \psi$  track-hit residuals in TIB as a function of the relative position of tracks from pp-collisions on the modules along the local  $u$ -axis before (magenta) and after (black) parametrization of the module shapes. Each residual is weighted by  $\tan^2 \psi$  of the track. Right: Mean probability of cosmic ray track fits as a function of their distance of closest approach to the nominal beam line for the different approaches to parametrize the module shapes.

mentum tracks from the interaction region cross the strip modules under small angles relative to the module normal. Therefore sensor curvatures have only a small effect. This is different for cosmic ray tracks that cross the tracker with a large closest distance to the beam line,  $d_0$ . The larger  $d_0$ , the larger the average track angle from the module normal, leading to degraded fit results for the flat module assumption, as shown in Fig. 9 (right). If curvature parameters on sensor level are determined, the average fit probability is almost flat as a function of  $d_0$  up to 50 cm, thus improving substantially the consistency between tracks from the interaction point and cosmic rays.

#### 3.1.2.4 Control of global deformations and weak modes

Generally, by minimising the function defined in equation 1 not all possible distortions of the tracker can be resolved. The residuals can be insensitive to certain global deformations referred to as *weak* modes. Weak modes correspond to global transformations that (at least approximately) preserve the validity of the track model for the track sample under consideration, i.e. the transformed hits are to a high degree still consistent with a valid trajectory. Such systematic distortions have very little impact on the goodness-of-fit of the track. However, these distortions might significantly bias the track parameters and other quantities such as invariant masses.

Systematic momentum biases for tracks with high momenta have therefore been investigated

using events with a  $Z$  boson decaying into oppositely charged muons. For example, without using the virtual  $Z$ -mass information in the alignment, a large dependence of the position of the mass peak as a function of the pseudorapidity  $\eta$  of the decaying positively charged muon is observed (Fig. 10, left).




Figure 10: Left: The mean reconstructed  $Z \rightarrow \mu\mu$  mass is shown as a function of the rapidity of the positive-charged muon. Results are compared from simulated events with perfect alignment (black), and 2011 data aligned without using the  $Z$  mass constraint (red) and 2011 data or aligned using the  $Z$  mass constraint (blue). Central: module-by-module position difference between the re-aligned geometry on top of Twist deformation and Summer 2011 geometry (after subtraction of global movements and rotations). The line indicates the twist applied. Right: Track  $\chi^2$  using loosely selected isolated muons from a statistically independent sample from the one used in the alignment procedure ( $p_T > 5\text{GeV}$ ).

This dependence can be attributed to a *twist* of the whole tracker defined as  $\Delta\phi_i = c \times z_i$  where  $\Delta\phi_i$  is the change of the azimuthal angle of a module  $i$ . In contrast, using  $Z$  mass information in the alignment fit, the remaining spread of  $Z$ -mass peak values is almost as small as in the detector simulation with perfect alignment.

In order to study the sensitivity of the alignment to weak modes, the possible deformations of the geometry are parametrized in the cylindrical coordinates  $r, z$ , and  $\phi$  [15] and applied to 2011 geometry. Afterwards, to test the capabilities of the alignment to correct for the introduced misalignment, the alignment procedure is rerun and then the module-by-module position differences with respect to the summer 2011 geometry are determined, subtracting global movements and rotations of the whole tracker, as shown in Fig. 10 (central). In addition, track  $\chi^2$  distributions for collision tracks are also shown for the summer 2011 geometry, the misaligned and the re-aligned geometry. It can be noticed that the twist deformation is indeed weak modes for collision tracks, because after the introduced misalignment the  $\chi^2$  distribution for collision tracks remain basically unchanged (Fig. 10, right). However, the introduced twist deformation is fully corrected by the alignment procedure using the  $Z$  boson decay information.

### 3.1.3 Tracker hit reconstruction

The first step of the reconstruction is referred to as *local reconstruction*. It consists of the clustering the signals on nearby pixels or strips into *hits*, and then estimating their position and its uncertainty, measured in the local coordinate frame of the sensor. An orthogonal coordinate system ( $u, v$ ) is used, with axes orientated along/perpendicular to the pixels/strips in the plane of the sensor. Of the two coordinates,  $u$  is defined to be the one that is perpendicular to the  $B$ -field, (except in the case of the stereo strip sensors, where it is defined as the one that makes the largest angle to the  $B$ -field).

For the pixel tracker, clusters are simply formed from one or more neighbouring pixels (either



side-by-side or corner-adjacent), each of which must have a signal exceeding 3200 electrons equivalent charge (with this requirement being applied by the pixel readout electronics). In addition the total signal of the cluster must exceed 4000 electrons equivalent charge. During track finding, a simple algorithm is used to estimate the hit position. Illustrating this for the  $u$  coordinate, the hit is first projected onto the  $u$ -axis. If it is only one pixel wide, the hit position is taken to be that of the centre of this pixel, corrected for the Lorentz shift. If it is wider than this, then its position is given by:

$$u_{rec} = u_C + \frac{Q_{last}^u - Q_{first}^u}{2(Q_{last}^u + Q_{first}^u)} |W^u - W_{inner}^u| - \frac{L_u}{2} \quad (2)$$

where  $Q_{first}^u$  and  $Q_{last}^u$  are, respectively, the charges collected in the first and last pixels of the cluster,  $u_C$  is the geometrical centre of the cluster,  $L_u$  is the Lorentz shift in the  $u$  direction,  $W_{inner}^u$  is the ‘inner width’ of the cluster, which is the width excluding the first and last pixel, and  $W^u$  is the projected cluster width, defined as:

$$W^u = D |\tan(\alpha^u - \pi/2) + \tan\Theta_L^u| \quad (3)$$

where  $D$  is the sensor thickness. The variable  $\Theta_L^u$  the Lorentz angle and  $\alpha^u$  is the track angle with respect to the sensor plane, both in the  $u$  direction.

A more sophisticated (and slower) estimate of the pixel hit position is made, when making a final fit to all the hits on a track to estimate its trajectory parameters. This is known as the ‘template-based reconstruction algorithm’ [16]. A detailed pixel sensor simulation, Pixelav [17], is used to simulate the expected cluster shape, if a charged particle crosses a sensor at a given angle and a given position within a pixel. This is called a ‘template’. Templates are produced for a wide range of angles and positions within a pixel, with each template recording the expected cluster shapes after projection onto the  $u$  and  $v$ -axes. The shapes of the clusters observed in data are compared with these templates, to see which one matches best. The position of the hit can then be inferred from the position of the particle used to produce that template.

For strip tracker, hits are formed from neighbouring strips, each of which must have a signal exceeding twice the rms noise. Each hit must also contain at least one strip with a signal exceeding three times the rms noise. Furthermore, the hit’s total charge should exceed the ‘hit noise’, defined as  $\sigma_{hit} = \sqrt{\sum_i \sigma_i^2}$ , where  $\sigma_i$  is the rms noise of strip  $i$ , and the sum runs over all strips in the hit. The position of each hit is estimated from the charge-weighted average of its strip positions, corrected for Lorentz shift of charge in the  $B$ -field. The position uncertainty is parametrized as a quadratic function of the projected cluster width  $W$ , defined in Eqn. 3, expressed in units of the strip pitch.

In the subset of strip tracker layers, in which each module has a second sensor, making a stereo angle to the first, 3-D hit positions are reconstructed by combining the information from the two sensors. The resulting 3-D hit are known as *matched hits*. As the two sensors are a few millimetres apart, this combination requires an estimate of the track trajectory (taken to be a straight line from the centre of CMS, if no track candidate is available).

The efficiency of the pixel and strip tracker sensors is measured with data using a sample of good quality tracks. In each sensor that the track trajectory crosses, the fraction of the time that a hit is recorded reasonably close to the track trajectory is measured. In the pixel and strip sensors, the efficiency is typically  $\approx 99.7 - 99.8\%$ , once the  $\approx 2.3 - 2.4\%$  of modules known to be defective are excluded from the measurement.

The hit resolution in the barrel sensors has been measured with data using the track hit residuals, which are defined as the difference between the hit position and the track position. The track is deliberately reconstructed excluding the hit under study in order to avoid bias. In the pixel detector, the resolution is measured in the middle of the three barrel layers, by comparing the predicted position with the measured hit. At high momentum, where multiple scattering between the layers does not affect the measurement, the hit resolution in the direction transverse to the beam axis is  $10.4\ \mu\text{m}$ . In the direction parallel to the beam axis, the resolution varies strongly as a function of the polar angle of the track: it is  $45\ \mu\text{m}$  for tracks moving perpendicular to the beam axis and improves to  $20\ \mu\text{m}$  for some more inclined tracks. (Optimum resolution is reached for a polar angle of  $\pm 60^\circ$ ). This variation is explained by the fact that more inclined tracks spread charge across several pixels, which allows a more precise determination of the hit position. In the strip detector, the uncertainty on the track position due to multiple scattering is unfortunately much larger than the inherent resolution, (because of multiple scattering), so an individual hit residual is not sensitive to the resolution. However, the difference in the track's hit residuals between two closely spaced modules can be measured with much greater precision, (because any offset in the track's position caused by multiple scattering will be largely common to both modules). This difference is measured for pairs of overlapping modules from the same strip track layer, and is used to infer the resolution. The strip sensor resolution depends strongly on the size of the cluster and on the pitch of the sensor. For barrel modules, it is typically in the range  $15\text{--}30\ \mu\text{m}$ .

### 3.1.4 Track reconstruction

The CMS tracking software is known as the 'Combinatorial Track Finder' (CTF). The collection of reconstructed tracks is produced by multiple passes (iterations) of the CTF track reconstruction sequence, in a process called *iterative tracking*. The basic philosophy of iterative tracking is that the early iterations search for the tracks that are easiest to find (e.g. relatively high  $p_T$  tracks produced near the interaction region). After each iteration, hits associated with already found tracks are removed, so reducing the combinatorics, and thus allowing later iterations to search for more difficult classes of tracks (e.g. low  $p_T$  or highly displaced tracks). For the data taken until August 2011, the iterative tracking consisted of 6 iterations. Iteration 0 is the source of most tracks and is designed to reconstruct prompt tracks with  $p_T > 0.8\ \text{GeV}/c$  and which have three pixel hits. Iteration 1 is used to recover prompt tracks which only have two pixel hits or slightly lower  $p_T$ . Iteration 2 is configured to find low  $p_T$  prompt tracks. Iterations 3–5 are intended to find tracks which originate outside the beamspot and to recover tracks not found by the previous iterations. At the beginning of each iteration, hits associated with *highPurity* tracks (defined later) found in previous iterations are masked off.

Each iteration proceeds in four steps:

- The *seed generation* provides initial track candidates found using only a few (2 or 3) hits. A seed defines the initial estimate of the trajectory parameters and their uncertainties.
- The *track finding* is based on a global Kalman filter [18]. It extrapolates the seed trajectories along the expected flight path of a charged particle, searching for additional hits that can be assigned to the track candidate.
- The *track fitting* module is used to provide the best possible estimate of the parameters of each trajectory by means of a Kalman filter and smoother.
- The *track selection* sets quality flags and discards tracks that fail certain criteria.

The main differences between the 6 iterations lie in the configuration of the seed generation



and final track selection steps.

#### 3.1.4.1 Track seeding

The trajectory seeds define the starting trajectory parameters and uncertainties of potential tracks. In the quasi-uniform magnetic field present in the Tracker, charged particles follow helices and therefore five parameters (including the trajectory curvature) are needed to define a starting trajectory. To obtain these five parameters requires hits in at least 3 tracker layers, or in 2 layers if a (loose) beam constraint is also used.

To limit the number of hit combinations, seeds are required to be compatible with the expected trajectory of a track with transverse momentum  $p_T$  exceeding a specified cut and transverse (longitudinal) impact parameter  $d_0$  ( $z_0$ ) less than a specified cut. The impact parameter here is usually calculated with respect to the reconstructed beam-spot position. However, it can also be calculated with respect to the reconstructed primary vertex, obtained with a 'pixel-only' track and vertex reconstruction algorithm, which because it uses only hits in the pixel detector is extremely fast. The beam-spot and primary vertex reconstruction algorithms are described in Sect. 3.1.4.7.

Table 3 shows the seeding layers and cuts used for each of the 6 tracking iterations. The 0<sup>th</sup> iteration is done using seeds formed from three pixel hits. The high 3-D hit resolution and low occupancy of the pixel tracker make it extremely easy to reconstruct tracks starting from these seeds. The 1<sup>st</sup> iteration uses pairs of hits in the pixel detector or strip tracker TID. It recovers tracks which, for any reason, failed to produce hits in three pixel layers, whilst by using the TID, it extends the coverage in the forward region. The 2<sup>nd</sup> iteration again uses three pixel hits, but now with a very low  $p_T$  cut. The 4<sup>th</sup> iteration uses triplets of hits containing at least one pixel and one strip tracker hit. It uses looser impact parameter cuts than the previous iterations, so allowing it to find non-prompt tracks from strange hadron decays, photon conversions and nuclear interactions. The final two iterations use 3-D 'matched hits' in pairs of those strip tracker layers that contain stereo sensors. They also use very loose impact parameters cuts, so allowing them to find highly displaced tracks.

Table 3: The configuration of the track seeding for each of the six iterative tracking steps. Shown are the layers used to seed the tracks, and the minimum  $p_T$  and maximum impact parameter requirements applied to the seeds. (The symbol '\*' indicates that the longitudinal impact parameter is calculated with respect to a pixel vertex, instead of the beamspot.)

Iteration	Seeding Layers	$p_T$ cut (GeV/c)	$d_0$ cut (cm)	$z_0$ cut
0	pixel triplets	0.8	0.2	$3\sigma$
1	mixed pairs with vertex	0.6	0.2	0.2 cm*
2	pixel triplets	0.075	0.2	$3.3\sigma$
3	mixed triplets	0.35	1.2	10.0 cm
4	TIB 1+2 & TID/TEC ring 1+2	0.5	2.0	10.0 cm
5	TOB 1+2 & TEC ring 5	0.6	5.0	30.0 cm

#### 3.1.4.2 Track finding

The track finding module of the CTF algorithm is based on the Kalman filter method. The filter begins with a coarse estimate of the track parameters provided by the trajectory seed, and then builds track candidates by adding hits from successive detection layers one by one. The

information provided at each layer includes the location and uncertainty of any found hit as well as the amount of material crossed, which is used to estimate the uncertainty arising from multiple Coulomb scattering. The Kalman filter method is implemented in four steps, listed below:

The first step, navigation, uses the parameters of the track candidate, evaluated at the current layer, to determine which adjacent layer(s) of the tracking detector, are intersected by the extrapolated trajectory, allowing for the current uncertainty on that trajectory. The navigation service uses a fast *analytical propagator* to find the layers. The analytical propagator assumes a uniform magnetic field and does not include the effects of multiple Coulomb scattering or energy loss. With these assumptions, the track trajectory is a perfect helix, so the propagator can extrapolate the trajectory from one layer to the next using analytical calculations.

The second step is to search in the layers returned by the navigation step for all sensors that the extrapolated track trajectory may cross, allowing for the statistical uncertainty on the trajectory.

The third step examines the hits collected from these compatible sensors. The hit positions and uncertainties are refined using the estimated track direction on the sensor surface. A  $\chi^2$  test is used to check which of the hits are compatible with the extrapolated track trajectory. The current (configurable) requirement is  $\chi^2 < 30$ . The  $\chi^2$  calculation takes into account both the hit and trajectory uncertainties. One can optionally also add an *invalid hit* to the collection of compatible hits, whose purpose is to allow for the possibility that the particle failed to produce a hit in that layer, for example, due to detector inefficiency. In addition, where two sensors within the same layer overlap slightly, a particle passing through the overlap region may produce two hits within that layer. The software allows both hits to be assigned to the track candidate.

The fourth and last step is to update the trajectory state. New track candidates are formed from each of the original ones, by adding to them exactly one of the compatible hits found in the third step (where this hit may be the invalid hit). The candidate's trajectory parameters are then updated at the new detector surface, by combining the information from the hit with the extrapolated track trajectory of the original candidate.

For the second, third and fourth steps above, a more accurate *material propagator* is used when extrapolating the track trajectory, which includes the effect of material in Tracker. This differs from the simple analytical propagator, in that it inflates the uncertainty on the trajectory parameters according to the predicted rms scattering angle in the Tracker material. It also adjusts the momentum of the trajectory according to the mean energy loss predicted by the Bethe-Bloch equation. Since all detector material is assumed to be concentrated in the detector layers, the track propagates along a helix between them, so allowing the material propagator to extrapolate the track analytically.

All resulting track-candidates found at each layer are then propagated to the next compatible layer(s) and the procedure is repeated until a 'stopping condition' is satisfied. However, to avoid an exponential increase of the number of candidates, only a limited number (default is 5) of candidates are retained at each step, with the best chosen based on the normalized  $\chi^2$  plus a bonus for each valid hit and a penalty for each invalid hit. The standard stopping conditions are if a track reaches the end of the Tracker detector or contains too many invalid hits, or if its  $p_T$  has dropped below a user specified value.

After a track candidate has completed the outward search for hits, an inward search for hits is begun. This is started by taking all of the hits assigned to the track, excluding those belonging to the track seed, and using them to fit the track trajectory. Then, following the steps above, this trajectory is propagated inwards through the seeding layers and then further inwards until

the inner edge of the Tracker is reached or too many invalid hits are found. There are two main reasons for this inward hit search: firstly, additional hits may be found in the seeding layers (for example, from overlapping sensors); secondly, hits may be found in layers closer to the interaction region than the seeding layers.

The track of a single charged particle may be reconstructed more than once, either starting from different seeds or when a given seed develops into more than one track candidate. To remedy this feature, a “trajectory cleaner” is applied after all the track candidates in a given iteration have been found. The trajectory cleaner calculates the fraction of shared hits between two track candidates:  $f_{shared} = \frac{N_{shared}^{hits}}{\min(N_1^{hits}, N_2^{hits})}$  where  $N_1^{hits}$  ( $N_2^{hits}$ ) is the number of hits in the first (second) track-candidate. If this fraction exceeds the (configurable) set value of 19%, the trajectory cleaner removes the track with the fewest hits; if both tracks have the same number of hits, the track with the largest  $\chi^2$  value is discarded. The procedure is repeated iteratively on all pairs of track candidates. The same algorithm is applied when tracks from the six iterations are combined into a single track collection.

The requirements applied during the track finding stage are shown in Table 4 for each tracking iteration. The minimum  $p_T$  requirements shown here have very little effect, as they are weaker than those applied to the seeds shown in Table 3. Since the later iterations do not have strong requirements that the tracks originate from the beamspot, the probability of random hits forming tracks increases, which leads to a higher fake rate and greater CPU time. To compensate for this tendency, the minimum hits and maximum lost hits requirements are tightened in the later iterations.

Table 4: The six step iterative tracking configuration gives the cuts applied during track finding. In addition to a minimum  $p_T$ ,  $N_{hit}$  is the minimum number of hits,  $N_{lost}$  is the maximum number of missing (invalid) hits, and  $N_{hit}^{rebuild}$  is the minimum number of hits needed in the in→out step to trigger an out→in step.

Iteration	$p_T$ cut (GeV/c)	$N_{hit}$	$N_{lost}$	$N_{hit}^{rebuild}$
0	0.3	3	1	5
1	0.3	3	1	5
2	0.1	3	1	5
3	0.1	4	0	5
4	0.1	7	0	5
5	0.1	7	0	4

### 3.1.4.3 Track fitting

For each trajectory, the track finding stage results in a collection of hits and an estimate of the track parameters. However, the full information about the trajectory is only available at the last hit of the trajectory and the estimate can be biased by constraints applied during the seeding stage. Therefore the trajectory is refitted using a Kalman filter and smoother.

The Kalman filter is initialized at the location of the innermost hit with the trajectory estimate obtained during seeding. The corresponding covariance matrix is scaled up by a large factor in order to avoid any bias. The fit then proceeds in an iterative way through the full list of hits, updating the track trajectory estimate with each hit in turn. For each valid hit, the hit’s position estimate is re-evaluated using the current values of the track parameters. This first filter is complemented with the smoothing stage: a second filter is initialized with the result of the first

one (except for the covariance matrix, which is scaled with a large factor) and is run backward toward the beam line. The track parameters at the surface associated with any of its hits, can then be obtained from the average of the track parameters of these two filters, evaluated on this same surface. (Since one filter uses information from all hits before that surface and the other uses information from all hits after the surface). Most usefully, one can thus obtain the track parameters evaluated at either the innermost or outermost hit on the track.

To obtain ultimate precision, this filtering and smoothing procedure uses a *Runge-Kutta propagator* to extrapolate the track trajectory from one hit to the next. This not only takes into account the effect of material, but is also able to accommodate an inhomogeneous magnetic field. The latter means that the particle may not move along a perfect helix, so its equations of motion in the magnetic field must be solved numerically. To do so, the Runge-Kutta propagator divides the distance to be extrapolated over into many small steps. It extrapolates the track trajectory over each of these steps in turn using a well-known mathematical technique for solving first order differential equations, called fourth-order Runge-Kutta method, which is accurate to fourth order in the step size. The optimal step size is chosen automatically, according to how non-linear the problem is.

Estimates of the track trajectory at any other points, such as the point of closest approach to the beam-line, can be obtained by extrapolating the trajectory evaluated at the nearest hit to that point. This extrapolation also uses the Runge-Kutta propagator.

#### 3.1.4.4 Track selection

The track finding algorithm described above yields quite a lot of fake tracks. Most physics analyses (and all tracking performance plots presented below) therefore use a subset of the tracks classified as *high purity*. The tracking software adds a flag to each track, indicating if it merits to belong to this class. Tracks must have a number of hits exceeding a certain cut, and not too many missing hits (in sensors where a hit would be expected). In addition, cuts on the track's  $\chi^2$  per degree of freedom and impact parameter significances are applied, with the these cuts being a function of the number of hits on the track, such that they are extremely loose for tracks with many hits.

#### 3.1.4.5 Reconstruction of electron tracks

Electrons, being charged particles, will be reconstructed by the standard tracking. However, due to their low mass, electrons often undergo bremsstrahlung in the detector material, which can lead to large energy losses and sharp changes in direction. The standard tracking algorithm uses a Gaussian approximation to model ionization energy loss and multiple Coulomb scattering. But it does not include the non-Gaussian tails that would be needed to model bremsstrahlung, and so suffers from relatively poor efficiency and resolution for electrons.

In consequence, CMS instead opts to reconstruct electron candidates, identified using energy deposits in the ECAL, using a dedicated Kalman filter tracking algorithm, that uses a 'Gaussian sum filter (GSF)' technique. This technique is described in detail in Ref. [19]. In essence, the fractional energy loss of electron, when it traverses material of given thickness, is predicted to have a distribution described by the Bethe-Heitler formula. This distribution is non-Gaussian, making in unsuitable for use in a conventional Kalman filter algorithm. The GSF technique solves this, by approximating the Bethe-Heitler distribution by the sum of several Gaussians.

### 3.1.4.6 Tracking performance

Figure 11 shows the tracking efficiency estimated using simulated events, each consisting of a single isolated particle. For muons of  $p_T > 0.7$  GeV/c, the efficiency is almost 100%. For pions, the efficiency is significantly lower, particularly in the endcap region. This is because of the high material budget of the tracker (see Sect. 3.1.1), which means that hadrons have a significant probability of undergoing a nuclear interaction before leaving the tracker volume. Figure 12 shows the tracking performance in simulated  $t\bar{t}$  events, generated with either no superimposed pile-up interactions or a mean of 8 superimposed pile-up interactions per event (typical of 2011 running conditions). The tracking efficiency is little affected by pile-up, and is very similar to that for single, isolated pions. The tracking algorithms thus perform well in a high occupancy environment. The fraction of fake tracks is  $\approx 5\%$ , peaking at  $|\eta| \approx 1.4$ , where the material budget is largest. It does rise with increasing pile-up.

Figure 13 shows the resolution in impact parameter and  $p_T$  single, isolated muons. (Similar results are obtained for hadrons in LHC physics events). For high  $p_T$  tracks (100 GeV/c), for which multiple scattering is of little importance, the resolution in the central region ( $|\eta| < 1$ ) is approximately 2% in  $p_T$  and 10  $\mu\text{m}$  (35  $\mu\text{m}$ ) in transverse (longitudinal) impact parameter.

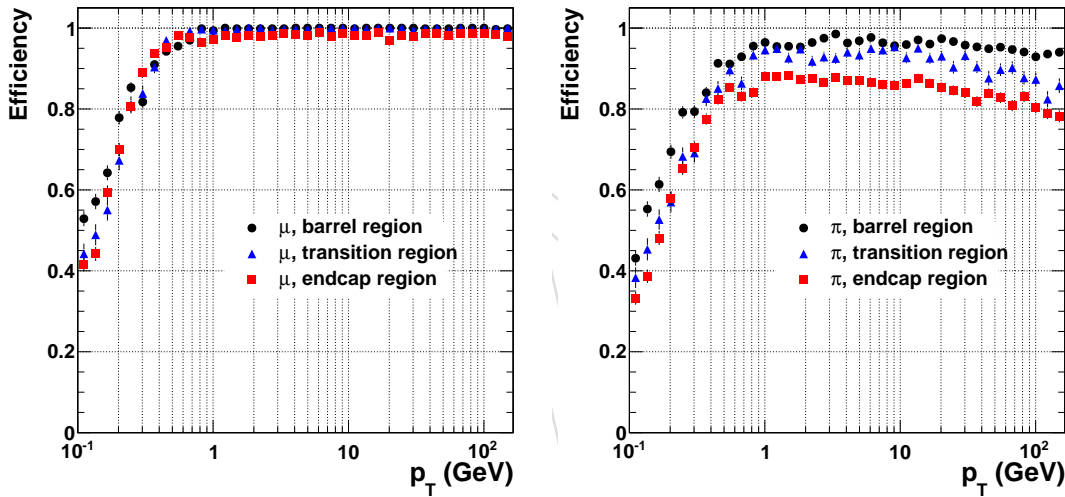


Figure 11: Track reconstruction efficiency for *single, isolated muons* (left) and *pions* (right) Results are shown as a function of  $p_T$  for the barrel, transition and endcap regions, that are defined by the pseudo-rapidity intervals [0-0.9], [0.9-1.4] and [1.4-2.5], respectively.

### 3.1.4.7 Primary vertex and beam-spot reconstruction

The goal of the primary vertex reconstruction [20], is to measure the location and uncertainty of all proton-proton interaction vertices in each event using the available reconstructed tracks. It begins by selecting good quality tracks produced near the beam-line. To ensure high reconstruction efficiency, even for minimum bias events, there is no requirement on the minimum allowed track  $p_T$ .

Track clustering (meaning assignment of tracks to one or more primary vertices) is performed with a deterministic annealing (DA) algorithm [21]. The  $z$ -coordinates of the points of closest approach of the tracks to the beamspot are referred to as  $z_i^T$ , and their associated uncertainties as  $\sigma_i$ . The tracks must be assigned to some unknown number of vertices at positions  $z_k^V$ . In the DA framework, assignments are soft. The  $p_{ik}$  are numbers between 0 and 1 and can be inter-






Figure 12: Tracking *efficiency* (left) and *fake rate* (right)  $t\bar{t}$ +jet(s) events simulated with and without superimposed pile-up collisions. The plots are produced using only tracks with  $p_T > 0.9$  GeV/c.

interpreted as the assignment probability of track  $i$  to vertex  $k$ . Hard assignment can be represented by values  $p_{ik}$  that equal 1 if track  $i$  is assigned to vertex  $k$  and 0 otherwise. In a large ensemble of possible assignments for a set of vertex positions. Postulating that a priori every possible configuration is equally likely, this is analogous to statistical mechanics when the vertex  $\chi^2$  takes the role of the energy. The most probable vertex positions at ‘temperature’  $T$  follow from minimizing what is called the free energy in statistical mechanics,

$$F = -T \sum_i p_i \log \sum_k \rho_k \exp \left( -\frac{1}{T} \frac{(z_i^T - z_k^V)^2}{\sigma_i^2} \right), \quad (4)$$

with respect to the vertex positions  $z_k^V$  and vertex weights  $\rho_k$ . The sums run over selected tracks ( $i$ ) and the set of vertices ( $k$ ) that are present at the current temperature. Tracks enter with constant weights,  $p_i$ , reflecting their compatibility with the beamspot. The number of vertices at constant temperature is fixed, but at the minimum of  $F$ , multiple vertices may assume identical positions. Rather than permanently using a very large number of vertices, the number of vertices coinciding at position  $k$  is represented by the weights  $\rho_k$ , such that the sum of all weights is constant at given temperature.

The temperature  $T$  controls the softness of the assignment and the resolution as effectively it scales the  $\sigma_i$  by  $\sqrt{T}$ . The clustering starts with a very high value of  $T$ , with all  $p_{ik}$  being equal and all tracks being compatible with a single vertex. The quantity  $F$  is minimised in a series of steps, with the temperature being reduced by a factor 0.6 at each step, and the numerical solution of each minimisation being initialised by the solution obtained at the previous step. The lower temperatures result in the minimisation finding additional vertices.

The DA algorithm is not robust against outliers, so in the final steps, tracks which are more than 4 standard deviations away from the nearest vertex are down-weighted in eqn. 4.

When  $T = 1$ , tracks are assigned to the vertex candidate for which they have the highest assignment probability, if it is greater than 0.5. Any vertex containing at least two tracks that are only compatible with that vertex are retained. The tracks from all retained vertices are then fit with an ‘adaptive vertex fitter’, described in Sect. ??, to compute the best estimate of

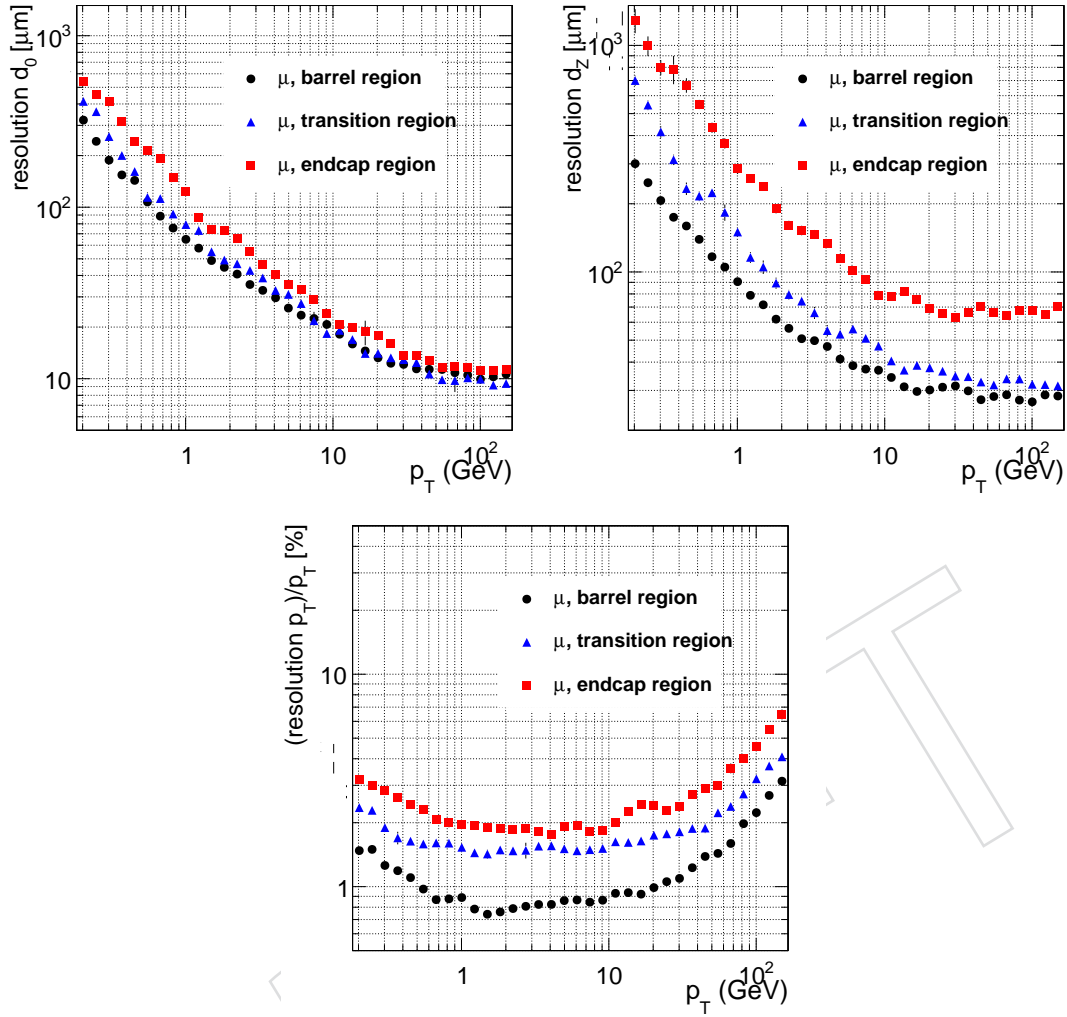


Figure 13: Resolution in  $d_0$ ,  $z_0$  and  $p_T$  for single, isolated muons in the barrel, transition and endcap regions, defined by the pseudo-rapidity intervals  $[0-0.9]$ ,  $[0.9-1.4]$  and  $[1.4-2.5]$ , respectively.

vertex parameters, including the position and covariance matrix, as well as the indicators of the success of the fit, such as the number of degrees freedom of the vertex and weights of the tracks in the vertex.

The primary vertex resolution depends strongly on the number of tracks used in fitting the vertex and the  $p_T$  of those tracks. It can be measured from data using the ‘split method’. In this method, the tracks assigned to a vertex are split equally into two different sets, in such a way that both sets have similar  $p_T$  distributions. These two different track sets are then fit independently with the adaptive vertex fitter. The distribution of the difference in the fitted vertex positions for a given number of tracks are fit with a single Gaussian distribution, and then divided by  $\sqrt{2}$ , to extract the resolution. Figure 14 shows the resolution in  $x$  (which is almost identical to that in  $y$ ) and in  $z$  as a function of the number of tracks in the vertex. Results are shown both for events selected with a minimum bias trigger and for events selected with a trigger requiring a jet of  $E_T > 20$  GeV. The resolution is better in the latter because of the higher track  $p_T$ , approaching  $10 \mu\text{m}$  in  $x$  and  $y$  and  $12 \mu\text{m}$  in  $z$ , for vertices with many tracks.






Figure 14: Primary vertex resolution in  $x$  (a) and  $z$  (b) as a function of the number of tracks, for two different kinds of events with track's of different average transverse momentum  $\overline{p_T}$ .

For very fast (but approximate) primary vertex reconstruction, ‘pixel tracks’ can be used. These fast primary vertices are used in certain track seeding steps, as described in Sect. 3.1.4.1, and are also used in the CMS trigger. Pixel tracks are reconstructed by simply searching for hits in three pixel layers that are compatible with a helical trajectory originating from near the nominal interaction point. Checking the compatibility of the three hits is done using simple geometrical calculations, without resorting to a Kalman filter. Such tracks can be found much more quickly than those found in the conventional way, although they suffer from poorer efficiency (since particles don’t always produced hits in three pixel layers), higher fake rates and worse resolution (particularly in  $p_T$ ). To estimate the primary vertex positions in a given event, the  $z$ -coordinate of the point of closest approach of each pixel track to the beam-line is calculated, and the tracks ordered according to this  $z$ -coordinate. A gap clustering algorithm is then employed: wherever two elements in this ordered list of  $z$ -coordinates are separated by distance exceeding a given cut value (that defaults to 1 mm), the tracks on either side of this point are assigned to separate vertices. The vertex positions are obtained by fitting the tracks assigned to each vertex with the adaptive vertex fitter, described in Sect. ??.

For events selected with a trigger requiring a jet of  $E_T > 20$  GeV, the pixel primary vertex resolution approaches  $35 \mu\text{m}$  in  $x$ ,  $y$  and  $z$ , for vertices with many pixel tracks.

The beam-spot represents the three-dimensional profile of the luminous region where the LHC beams collide at CMS. Its position and size is reconstructed by averaging over many events, in contrast to the primary vertex position, which is reconstructed event-by-event. It provides an estimate of the nominal primary vertex position, which is available before primary vertex reconstruction has been performed. It is used by the track seeding steps, as described in Sect. 3.1.4.1. The three-dimensional beam-spot position and width can be obtained from the distribution of reconstructed primary vertices. However, the position can also be obtained from a second method, which exploits a correlation between the transverse impact parameter ( $d_0$ ) and the azimuthal angle ( $\phi$ ) of tracks that exists if the beamspot is displaced from the expected position. To first order the  $d_0$  for tracks coming from the primary vertex can be parametrized by

$$d_0(\phi, z) = x_0 \cdot \sin \phi + \frac{dx}{dz} \cdot \sin \phi \cdot z - y_0 \cdot \cos \phi - \frac{dy}{dz} \cdot \cos \phi \cdot z, \quad (5)$$

where  $x_0$  and  $y_0$  are the position of the beam at  $z = z_0$ , and  $\frac{dx}{dz}$  and  $\frac{dy}{dz}$  are the  $x$  and  $y$  slopes of the beam. The beamspot fit [22] uses an iterative  $\chi^2$  fit to determine all four of these parameters

(although the  $z$ -coordinate of the centre of the beam-spot is determined from the distribution of primary vertices). With a sample of 1000 tracks, the position in the transverse plane can be determined with a statistical precision of  $\sim 5 \mu\text{m}$ . The position of the beam-spot is reconstructed using an average over all the events in a ‘luminosity section’ (which is defined as a period 23 s long). An average over more than one luminosity section (with a maximum of sixty) can be used if they all give consistent results.

### 3.2 Electrons/Photons

G. Franzoni

### 3.3 Muons

Muon detection is a powerful tool for recognizing signatures of interesting processes over the very high background rates at the LHC. This becomes more important as the luminosity increases. The CMS muon detector system has 3 functions: muon identification, momentum measurement, and triggering. Good muon momentum resolution is provided by the high-field solenoidal magnet and the flux-return yoke, which also serves as a hadron absorber that reinforces muon identification.

The CMS muon system (*ref to det picture*) is designed to reconstruct the momentum and charge of muons over the entire kinematic range of LHC collisions (Fig. 15). CMS uses 3 types of gaseous particle detectors for muon identification. Because of the central solenoidal magnet, it is natural to have a cylindrical barrel region and planar endcaps.




Figure 15: Invariant mass spectrum of dimuons in events collected with the loose double-muon trigger in 2010.

The basic detector process utilized in the CMS Muon System is gas ionization. For all 3 technologies—drift tubes, cathode-strip proportional planes, and resistive plates—the basic physical modules are called “chambers”. The chambers are independently-operating units, which are assembled into the overall muon detector system of CMS. The chambers form part of a spectrometer in which the analyzing magnet is the central solenoid together with the flux return yoke of CMS. To match the cylindrical geometry of the solenoid, the barrel region is instrumented with drift

tube chambers, and the 2 endcap regions with cathode strip chambers. Resistive plate chambers are interspersed in both the barrel and endcap regions. To utilize the magnet to measure the deflection of muons as they pass through its field, the muon chambers must detect the traversing track at several points along the track path. In the barrel region this means that chambers are positioned at several different values of the radial distance  $r$  from the beam line, and in the endcap region at several different values of distance along the beam direction  $z$ . A “station” is an assembly of chambers around a fixed  $r$  (in the barrel) or  $z$  (in the endcap) value. There are 4 stations in the barrel and in each endcap (*ref to det picture*), labelled MB1–MB4 and ME1/n–ME4/n, where integer  $n$  increases with the radial distance from the beam line.

The barrel drift tube (DT) chambers cover the pseudorapidity region  $|\eta| < 1.2$ . They are organized into 4 stations at different radii and interspersed between the yoke return plates. Each station consists of 8 layers of tubes measuring the position in the bending plane and 4 layers for the longitudinal plane.

In the endcap regions of CMS, where the muon rates and background levels are high and the magnetic field is strong and non-uniform, CMS uses cathode strip chambers (CSC). These chambers have a fast response time (because of the very short drift path), fine segmentation, and relative immunity to the non-uniformity of the field. The CSCs cover the  $|\eta|$  region from 0.9 to 2.4. Each endcap has 4 stations of chambers mounted on the faces of the endcap steel disks perpendicular to the beam. A CSC consists of 6 layers, each of which measures the muon position in 2 coordinates. The cathode strips run radially outward and provide a precision measurement in the  $r$ – $\phi$  bending plane. The wires provide a coarse measurement in the CMS radial direction. The CSCs in the innermost ring of station 1 have a structure different from those of the other rings in order to cope with specific geometrical requirements and the influence of the strong magnetic field in this region of the CMS detector.

These muon detector elements cover the full CMS pseudorapidity interval  $|\eta| < 2.4$  with no acceptance gaps ensuring good muon identification over a range corresponding to  $10^\circ < \theta < 170^\circ$ . In addition to these tracking detectors, CMS has added a complementary, dedicated triggering detector with excellent time resolution to reinforce the measurement of the correct beam-crossing time at the highest LHC luminosities. The resistive plate chambers (RPC) are located in both the barrel and endcap regions, and they can provide a fast, independent trigger with a looser  $p_T$  threshold over a large portion of the rapidity range ( $|\eta| < 1.6$ ). The RPCs are double-gap chambers, operated in avalanche mode to ensure reliable operation at high rates.

### 3.3.1 Local muon trigger

The local muon triggers provide the trigger segments for the Level-1 trigger from each barrel and endcap muon chamber. In the barrel, this task is performed by the DT local trigger (DTLT), and in the endcaps, by the CSC local trigger (CSCLT). The RPC trigger is not based on local trigger devices, as muon trigger candidates are constructed from the spatial and temporal coincidence of hits in the RPCs. The DT and CSC local trigger segments from each muon station are collected by the trigger Track Finder (TF), which combines them to form a muon track and assigns a transverse momentum value. At least 2 segments in 2 different stations are needed by the TF to construct a muon candidate.

The timing parameters of the LHC are such that proton bunches potentially cross inside CMS every 25 ns. The purpose of the trigger system is to find event candidates fulfilling a predefined set of criteria and to assign them to an appropriate bunch crossing number. Trigger synchronization of a subsystem must be achieved at 3 levels: intrachamber synchronization, chamber-to-chamber relative synchronization, and subsystem-to-subsystem synchronization. Although

each muon subsystem faced unique challenges due to differences in chamber design, trigger electronics design, and physical position on the CMS detector, the general synchronization procedures were similar.

The time of the hit left by the passage of a particle through the muon chamber with respect to the locally distributed clock signal depends on the muon time of flight from the interaction point to the chamber, individual chamber properties and/or their geometrical position, latency of trigger electronics and length of cables and fibers connecting the chamber electronics to the peripheral crates. These properties are specific to each chamber and were studied during cosmic data taking, before pp collisions were recorded at CMS. The synchronization with respect to the master LHC clock frequency (and hence with the rest of CMS) is achieved by moving the phase of the locally distributed BX with respect to the master LHC clock. The internal delay settings were further tuned to give optimal phase with respect to the LHC machine clock using high  $p_T$  muon data from first LHC collisions.

A dedicated DAQ stream was developed to collect at high rate a fraction of the CMS RAW data content consisting only of L1 trigger information, before any HLT processing. By comparing the distributions of L1 muon trigger times to the expected collision times, one can measure the L1 synchronization. After contamination from cosmic-ray muons and beam halo was subtracted, the fraction of pre- and post-fired events was on the order of 0.2% and 0.1%, respectively *need to check if this is true for DT*. These results exceeded the Physics TDR expectations of 99% in-time triggering [23].

The position and angle of the offline reconstructed track segments are compared with the corresponding information assigned by the DTLT to the trigger segments to determine the position and angular resolution of the DTLT primitives. The RMS of the position difference distribution is approximately 1 mm, and is the same for every station type. The RMS of the angular difference distribution is 4.8 mrad. The results are in agreement with previous measurements [24, 25], showing that the expected performance in terms of position and transverse momentum resolution at the output of the Level-1 trigger is achieved [26]. In a similar way, the position resolution of the CSCLT primitive is measured by comparing the position of the the CSC trigger primitive in the station to the position of the best-matched CSC track segment defined offline. The measured resolution is 3.2 mm, and 2.2 mm for ME1/1, which has narrower cathode strips. The angular resolution is much less critical in the endcaps, owing to their more complete geometrical coverage. For this reason the direction resolution of the CSC trigger primitives was not studied in detail.

Table 5: Average DTLT efficiencies for the different station types, for data and simulation. Results that include the correct BX identification (BXID) are also shown. Error values include both statistical and systematic uncertainties.

Station	DTLT Efficiency (%)		BXID Efficiency (%)	
	Data	Simulation	Data	Simulation
MB1	$96.2 \pm 0.1$	$97.9 \pm 0.9$	$94.5 \pm 0.9$	$96.4 \pm 0.9$
MB2	$95.7 \pm 0.8$	$98.0 \pm 0.8$	$94.0 \pm 0.8$	$96.7 \pm 0.7$
MB3	$95.8 \pm 0.9$	$98.3 \pm 0.8$	$93.8 \pm 0.9$	$96.9 \pm 0.8$
MB4	$95.0 \pm 0.1$	$97.1 \pm 0.9$	$93.0 \pm 0.9$	$95.6 \pm 0.9$

To measure the DTLT efficiency, selected events are required to be triggered by the RPC system, without any requirement on the presence of the DT trigger, which could otherwise bias the measurement. The presence of a well reconstructed reconstructed muon track in the DT accep-

tance is then required. The contamination from cosmic rays is removed by selecting only tracks with small impact parameter. The track transverse momentum is required to be  $p_T > 7$  GeV/c, to allow the particle to reach the outer station of the muon barrel. The average DTLT efficiency is shown in Table 5, for data and simulation, for the 4 barrel stations. The overall systematic uncertainty, which dominates the measurement, is estimated from the observed spread of the measured DTLT efficiencies over the various stations, once the stations with known hardware problems are removed. The lower efficiency measured in the data with respect to the simulation is partially due to small differences in the timing of the muon stations from ideal conditions, and is a subject of further investigation during the 2011 data-taking campaign *to be checked if any news*. If in the DTLT efficiency definition the trigger primitive is also required to correctly assign the BX at which the muon candidate is produced, the average efficiency decreases to 93.8%, which is more than 1% better than the design (Level-1 Trigger TDR) performance [26]. Results for the DTLT efficiency including the correct BX assignment are shown in Table 5.

Table 6: Average CLCT efficiencies per station: the uncertainties are statistical and systematic, respectively.

Station	Cathode Trigger Primitive Efficiency (%)			
	Single Track Matching Method		Tag-and-Probe Method	
	Data	Simulation	Data	Simulation
ME1	$97.9 \pm 0.1 \pm 0.7$	$99.0 \pm 0.1 \pm 1.3$	$98.7 \pm 0.9 \pm 0.8$	$97.2 \pm 0.1 \pm 1.3$
ME2	$97.0 \pm 0.1 \pm 0.4$	$96.7 \pm 0.1 \pm 0.4$	$95.6 \pm 0.9 \pm 1.0$	$94.2 \pm 0.2 \pm 1.4$
ME3	-	-	$96.0 \pm 0.9 \pm 1.1$	$92.5 \pm 0.2 \pm 1.9$
ME4	-	-	$94.5 \pm 1.6 \pm 1.4$	$89.8 \pm 0.3 \pm 2.1$

Two methods were used to measure the CSCLT efficiency: the “single track matching” and the “tag-and-probe” methods. Both methods use tracks reconstructed using only tracker detector information (the so-called “tracker tracks”), to perform an unbiased measurement. In the “single track matching” method, a high quality tracker track whose extrapolation crosses a CSC station (tag station) is selected, and the CSCLT efficiency is measured in another station (probe station) that is in front of the tag station. In the tag-and-probe method, muon candidates from  $J/\psi$  and  $Z^0$  decays are used. Of the two muon candidates from a resonance decay, one is used as a tagging track and the other acts as a probe.

The CSCLT efficiency is measured for both data and simulated events. The single track matching method is applied to minimum bias data, whereas the tag-and-probe method is applied to the  $J/\psi$  and  $Z^0$  samples. The average efficiencies for stations ME1 and ME2 are reported in Table 6. The uncertainties are dominated by systematic effects, related to the selection criteria applied to the tracks and the trigger segments. Only in the case of the tag-and-probe measurement of data are the statistical and systematic uncertainties of similar size.

### 3.3.2 Local muon reconstruction

The global reconstruction of muons relies on the local reconstruction of objects inside the individual muon chambers. The 3 muon detector systems (DT, CSC, and RPC) use different techniques to register and reconstruct signals originating from charged particles penetrating them. Still, in all cases the basic objects are “reconstructed hits” (i.e., 2- or 3-dimensional spatial points with assigned uncertainties) and segments, obtained by fitting straight lines to the reconstructed hits.



The track segment reconstruction in the barrel DT chambers proceeds in 2 steps: first, a hit reconstruction consisting of deriving spatial points from the TDC time measurements; second, a linear fit of these points in the 2 projections of a chamber (8  $\phi$ -layers and 4  $\theta$ -layers), to perform a local pattern recognition and obtain reconstructed segments. The first step starts with the calibration of the TDC output to get the real drift times of the ionized charge within the tube. Since a DT cell is 42 mm wide and has a central wire, the maximum drift distance is 21 mm. Then, multiplying the drift times for a known drift velocity, a pair of space points (rechits) are obtained, left and right, at equal distance from the wire. These are the input to the linear fit that attempts to associate the majority of rechits to a segment. Hits that are inconsistent with the fit, yielding high segment  $\chi^2$  values, are discarded. At least 3 rechits from different layers are required to build a segment.

In the endcap CSCs, the rechit reconstruction is based on information from the strips (local  $x$  or  $\phi$  coordinates) and wires (local  $y$  coordinate). A rechit is built only if signals from both strips and wires are present in a given layer. The strip width varies between 0.35 and 1.6 cm for different chamber sizes and locations, and a typical muon signal is contained in 3 to 6 strips. The charge distribution of the strip signals is well described by a Gatti function ([27], [28], [29]), which is the basis of the the local  $x$  coordinate reconstruction. The wire signals are grouped in wire groups with widths between 2 and 5 cm and typically only 1 wire group is fired. The uncertainties on the position measurements are extracted from studies performed from measurements using both cosmic and collision data. The proper error correlation matrix is assigned to a rechit (strips and wires are not perpendicular). Segments are built from the available rechits in the 6 layers of the chambers and and rechits from at least 3 layers are required to build a segment. Rechits that are rejected by the fit have positions typically distorted by the presence of delta electrons.

The RPC local reconstruction has as input the strips that were fired on a given event. The strips that are next to each other are grouped into what is called a “strip cluster”, and the average position of the strips that form the cluster is what constitutes the reconstructed hit of a given RPC detector. The uncertainty on the measurement follows the standard deviation of a uniform distribution (i.e., simply the size of the cluster along each direction divided by  $\sqrt{12}$ ).

Hit resolution is determined from the distribution of hit residuals with respect to the muon trajectory. This is possible in the DTs and CSCs with no need of an external reference by using the track stubs (“segments”) reconstructed with a straight line fit of the hits in the different measurement layers. Therefore, the relative alignment of chambers does not affect the result. The residual of hits with respect to the reconstructed segment is a biased estimator of the resolution because the hit under study contributes to the segment fit if all available hits are included in the segment, or because of the uncertainty in segment extrapolation or interpolation if the fit is performed after removing the hit under study. In either case, the bias can be corrected for by using the statistical relationship between the width of the residual distribution for a given layer and the actual resolution. Monte Carlo studies have shown these corrections allow the true resolution to be obtained with good accuracy. Residuals in RPC chambers, which provide a single measurement of the trajectory, are defined by extrapolating the segment of the closest DT or CSC chamber.

The single-hit DT resolution is summarized in Table 7. The resolution is approximately the same for positive and negative pseudorapidity given the geometric symmetry of the system. In Wheel 0, the resolution is the same for the  $r$ - $\phi$  and  $r$ - $z$  layers. The resolution varies moving towards external wheels due to the effect of the larger longitudinal incidence angle ( $\theta$ ) of muons on these chambers. For  $r$ - $z$  SLs,  $\theta$  is the angle in the measurement plane; therefore the

resolution is significantly degraded in external wheels, due to the increasing deviation from linearity of the space-time relationship with larger angles of incidence of the particles. For  $r$ - $\phi$  layers,  $\theta$  is the angle in the plane orthogonal to the measurement plane; the larger angle in external wheels results in longer paths inside the cells that increase the primary ionization statistics, causing a slight improvement in the  $r$ - $\phi$  resolution. The poorer resolution of the  $r$ - $\phi$  layers in MB4 compared to MB1–3 is due to the fact that in this station, where a measurement of the orthogonal coordinate is missing, it is not possible to correct for the actual muon time of flight and signal propagation time along the wire. This correction can be applied at a later stage, during the fit of a muon track using all stations.

Table 7: Single-hit DT resolution for  $r$ - $\phi$  and  $r$ - $z$  layers expressed in  $\mu\text{m}$ , averaged over all sectors in each barrel wheel and station.

SL Type	Station	W-2	W-1	W0	W1	W2
$r$ - $\phi$	MB1	$220 \pm 7$	$249 \pm 16$	$266 \pm 8$	$243 \pm 9$	$217 \pm 7$
	MB2	$229 \pm 12$	$255 \pm 10$	$276 \pm 9$	$252 \pm 6$	$223 \pm 15$
	MB3	$229 \pm 11$	$256 \pm 11$	$275 \pm 10$	$255 \pm 10$	$231 \pm 8$
	MB4	$287 \pm 21$	$325 \pm 26$	$349 \pm 23$	$315 \pm 21$	$284 \pm 19$
$r$ - $z$	MB1	$772 \pm 32$	$432 \pm 25$	$266 \pm 16$	$419 \pm 18$	$785 \pm 33$
	MB2	$627 \pm 26$	$425 \pm 19$	$280 \pm 10$	$419 \pm 16$	$651 \pm 43$
	MB3	$538 \pm 30$	$372 \pm 13$	$278 \pm 14$	$368 \pm 18$	$538 \pm 27$

The spatial resolution of the segment fitted in the whole chamber, obtained as  $\sigma/\sqrt{N}$ , where  $\sigma$  is the single-hit resolution and  $N$  the number of layers included in the fit, is reported in Table 8.

Table 8: DT chamber resolution in the  $r$ - $\phi$  and  $r$ - $z$  projections, expressed in  $\mu\text{m}$ , averaged over all sectors in each barrel wheel and station.

SL Type	Station	W-2	W-1	W0	W1	W2
$r$ - $\phi$	MB1	$78 \pm 2$	$88 \pm 6$	$94 \pm 3$	$86 \pm 3$	$77 \pm 2$
	MB2	$81 \pm 4$	$90 \pm 3$	$98 \pm 3$	$89 \pm 2$	$79 \pm 5$
	MB3	$81 \pm 4$	$90 \pm 4$	$97 \pm 4$	$90 \pm 4$	$82 \pm 3$
	MB4	$101 \pm 7$	$115 \pm 9$	$123 \pm 8$	$111 \pm 7$	$100 \pm 7$
$r$ - $z$	MB1	$386 \pm 16$	$216 \pm 12$	$133 \pm 8$	$209 \pm 9$	$393 \pm 17$
	MB2	$314 \pm 13$	$212 \pm 9$	$140 \pm 5$	$210 \pm 8$	$325 \pm 22$
	MB3	$269 \pm 15$	$186 \pm 7$	$139 \pm 7$	$184 \pm 9$	$269 \pm 14$

The CSC spatial resolution is determined by the design parameters of the chambers, as well as certain characteristics of each muon track and their reconstruction. The measured chamber resolutions obtained using different samples and taking into account different geometrical characteristics of the chambers are summarized in Table 9.

The resolutions obtained from cosmic muons are found to be slightly worse than those measured from collision data. This is expected due to several reasons: cosmic muons arrive uniformly distributed in time, which worsens the hit resolution; they have a larger variation of angles of incidence; and there are backgrounds that can be removed from the collision events by requiring the presence of well reconstructed muons with  $p_T > 20 \text{ GeV}/c$  that cannot be removed from cosmic muon data. Reasonable agreement with simulated events is found for pp collision data. There are known differences for the ME1/1 chambers (in parentheses in Table 9): while the HV in the data was lowered by 4%, the MC still uses the design HV setting<sup>1</sup>, which

<sup>1</sup>During 2010 running, the HV for the ME1/1 chambers was set lower than the design value because the design



Table 9: Average CSC position resolutions ( $\mu\text{m}$ ).

Run type/period	Chamber type							
	ME1/1	ME1/2	ME1/3	ME2/1	ME2/2	ME3/1	ME3/2	ME4/1
Cosmics 2009	70	109	129	144	158	143	197	208
Cosmics 2010	70	109	130	144	157	143	193	196
Collisions 2010	58	92	103	126	132	126	136	131
pp MC 2010	(37)	80	110	121	150	115	151	114

leads to the observed better resolution in the simulation. Further smaller differences seen between data and simulation are understood to be due to differences in simulating the deposited charge, which will be addressed in future work. All measured resolutions are close to and most even exceed the requirements noted in the CMS Muon TDR [30], which called for 75  $\mu\text{m}$  for the ME1/1 and ME1/2 chambers and 150  $\mu\text{m}$  for the remaining chambers.

In the CMS muon system, RPCs are used as trigger detectors; nevertheless, hits are provided for reconstruction and muon identification. To measure the resolution of the RPC system, DT and CSC segment extrapolation was performed and compared with the reconstructed RPC hit (the average position of the strips fired when a muon passes through a given RPC chamber). The Gaussian fit results for all the different RPC residual distributions in CMS (Table 10) should be considered as the definitive resolution measurements for the RPC system. As expected, a clear correlation can be seen between RPC resolution and strip pitch.

In addition to a measurement of the track position and direction, the DTs and CSCs provide a measurement of the arrival time of a muon in a chamber, together with the RPCs. In a sample of prompt muon tracks the overall segment time resolution was determined to be better than 2.4 ns in DT and better than 3 ns in CSC and RPC.

The hit reconstruction efficiency is measured using reconstructed local segments. However, to reduce possible biases, loose selection criteria were applied to the reconstructed segments in order to discard poor quality segments originating from fakes and background particles.

Figure 16, left, shows both hit reconstruction and hit association efficiency as functions of the position in a DT cell. Apart from the known inefficiency induced by the cathode at the edges of the cell [31], the hit reconstruction efficiency is everywhere  $\geq 99\%$ . The hit association efficiency is, as expected, up to 2% lower, as it depends on the details of the calibration and contributions from  $\delta$ -rays. Figure 16, right, shows the hit reconstruction and the hit association efficiencies as functions of position in the layer for a subset of DT MB1 chambers. The efficiency is constant along the layer and the cell structure is clearly visible. Overall, the hit reconstruction efficiency in the barrel DT system is on average  $\approx 98\%$ , whereas the association efficiency is  $\approx 96\%$ .

Figure 17 shows the rechit efficiency in the endcap CSCs of station 2, ring 2, for all layers as a function of the local  $y$  coordinate (left) and the strip  $\phi$  angle (right). The “dead” chamber regions located between the high voltage segments are clearly visible on the left plot. A slight inefficiency is observed at the boundaries between consecutive cathode readout boards (CFEB) in the  $\phi$  efficiency plot (right). The rechit efficiency in the “active” CSC regions is well above 99.5%. The association efficiency is  $97.1 \pm 0.1\%$  for the 4 central (out of 6) layers and there is a small bias ( $< 1\%$ ) in the 2 outer layers related to the procedure of removing “bad” hits inside a segment.

The barrel and endcap RPC systems are mainly used as trigger detectors; however they do

---

resolution could be obtained at the lower value, which increases the lifetime of the chambers. If necessary, the HV can be increased and the spatial resolution improved further.

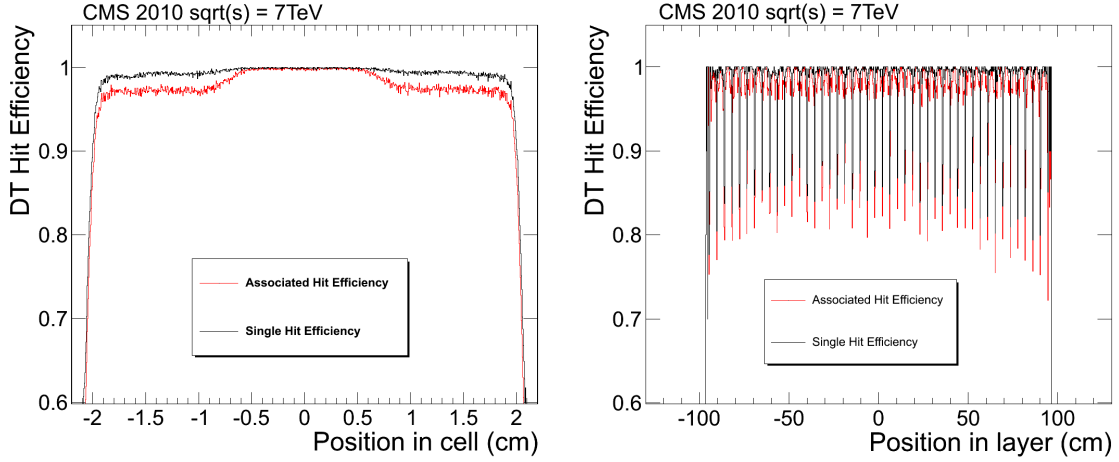


Figure 16: Left: Hit reconstruction (black) and hit association (red) efficiency as a function of the track position in a DT cell. Right: Hit reconstruction (black) and hit association (red) efficiency as a function of the track position in a DT MB1 layer .




Figure 17: CSC rechit reconstruction efficiency based on the segment propagation as functions of the local  $y$  (wire groups) coordinates (left) and strip  $\phi$  coordinates (right).

contribute to the muon reconstruction by providing additional position and time information to that provided by the DT and CSC systems in the barrel and endcap regions. Every RPC is located close to a DT or CSC and therefore the extrapolation of a segment reconstructed by the latter should point to a specific RPC strip and to a particular location within the strip. In a sense, an RPC can be considered as an additional DT or CSC layer. This allows the use of reconstructed DT and CSC segments as probes for determining the RPC efficiency.

The expected rechit efficiency of the RPC system is 95%. The measured efficiency is shown in Fig. 18, separately for the barrel (left) and endcap (right) RPCs. As can be seen, the efficiencies in the barrel and endcap chambers are comparable and close to expectations. The tail of lower efficiency chambers is a result of RPCs affected by electronics problems (e.g., a few dead channels) or not operating at the optimal voltage. The percentages of RPCs with efficiencies below 80% in the barrel and endcap are 1.2% and 1.3%, respectively. For the 2011 data-taking period, a set of calibration runs has been taken to tune the operating voltage chamber by chamber.

The segment reconstruction efficiency measurement was performed using the tag-and-probe method applied to well-identified muons from  $J/\psi$  and  $Z^0$  decays selected from 2010 collision data at  $\sqrt{s} = 7$  TeV. Corresponding samples of simulated events were used to compare the observed efficiency distributions with the expected one.




Figure 18: RPC efficiency distribution per chamber for the barrel (left) and endcap (right). Chambers with known problems are excluded from the plot.

For the DTs, the passing probes are those that were matched by the global track reconstruction to a segment reconstructed within the chamber. For the CSCs, the passing probes are required to match a local CSC segment within a distance of 40 cm, which is a loose requirement but further reduces the small background from cosmic muons or other muons and jets present in the event.

The selected probe tracks were propagated to the muon stations, starting from their point of closest approach to the interaction point. The propagation procedure allows the position of the track to be determined at any crossed surface. Uncertainties due to multiple scattering on the extrapolated position match the MC expectations.

The presence of reconstructed segments was checked for each individual chamber crossed by the probe tracks. To reduce the apparent loss of efficiency owing to propagation errors, the intersections between probe track and chamber were required to be away from the chamber edges, reduced by the error on the position of the intersection itself. The residual effects of propagation errors are included in the measured efficiency.

Owing to energy losses in the traversed material there is a minimum momentum (or  $p_T$  at a given  $\eta$ ) threshold for muons to reach the muon detector. The  $p_T$  threshold is  $\approx 1$  GeV/ $c$  for the forward region and increases to  $\geq 4$  GeV/ $c$  in the central region. To reduce the effect of multiple scattering on the efficiency measurement and to assure the muon has the energy to further penetrate all the muon stations a requirement on the minimal  $p_T$  ( $p$ ) of the track probes is imposed. In the following, a selection of  $p_T > 10$  GeV/ $c$  in the barrel and  $p > 15$  GeV/ $c$  in the endcap was applied (except for the  $p_T$  dependent measurement).

Figure 19 shows the segment efficiency as a function of  $p_T$  of the probe track, for the 4 barrel DT stations.

Figure 20 shows the segment efficiency as functions of  $\eta$ , and  $p_T$  of the probe track, for 2 endcap CSC stations and the comparisons with simulation.

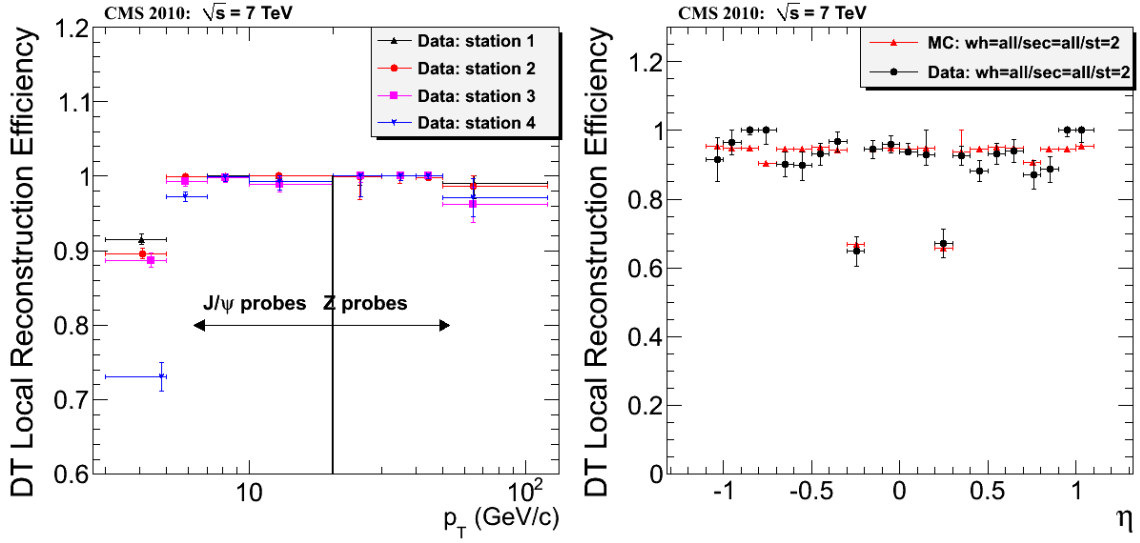


Figure 19: Segment reconstruction efficiency as a function of transverse momentum in the 4 barrel DT stations (left); the vertical line separates the ranges covered by probes originating from  $J/\psi$  and  $Z^0$  decays; and as a function of  $\eta$  in station 2, compared to simulated data (right); no cut around chamber edges is applied for this plot.

The overall segment reconstruction efficiencies measured in the endcap (CSC) and barrel (DT) muon systems are summarized in Table 11. Here the systematic uncertainties are obtained by estimating the impact of the multiple-scattering effects on the probes.

In summary, the reconstructed segment efficiency determined using the tag-and-probe method with real data is at the level of 95%–99% in all muon system stations with a systematic uncertainty of less than 2%. The dead zones in the muon chambers and the non-operating chambers/regions are the main reasons for segment inefficiencies in the muon chambers. We see an overall good agreement between data and Monte Carlo simulation within the uncertainties.

### 3.3.3 Global muon reconstruction

In the CMS reconstruction for pp collisions [6, 32], tracks are first reconstructed independently in the inner tracker (*tracker track*) and in the muon system (*standalone-muon track*). Based on these objects, two muon reconstruction approaches are used:

- *Global Muon reconstruction (outside-in)*. For each standalone-muon track, a matching tracker track is found by comparing parameters of the two tracks propagated onto a common surface, and a *global-muon track* is fitted combining hits from the tracker track and standalone-muon track, using the Kalman-filter technique [18]. At large transverse momenta,  $p_T \gtrsim 200$  GeV/c, the global-muon fit can improve the momentum resolution compared to the tracker-only fit [6, 32].
- *Tracker Muon reconstruction (inside-out)*. In this approach, all tracker tracks with  $p_T > 0.5$  GeV/c and the total momentum  $p > 2.5$  GeV/c are considered as possible muon candidates and are extrapolated to the muon system taking into account the magnetic field, the average expected energy losses, and multiple scattering in the detector material. If at least one muon segment (i.e., a short track stub made of DT or CSC hits) matches the extrapolated track, the corresponding tracker track qualifies as a Tracker Muon. Track-to-segment matching is performed in a local (chamber) coordinate system, where local  $x$  is the best-measured coordinate (in the  $r$ - $\phi$  plane)

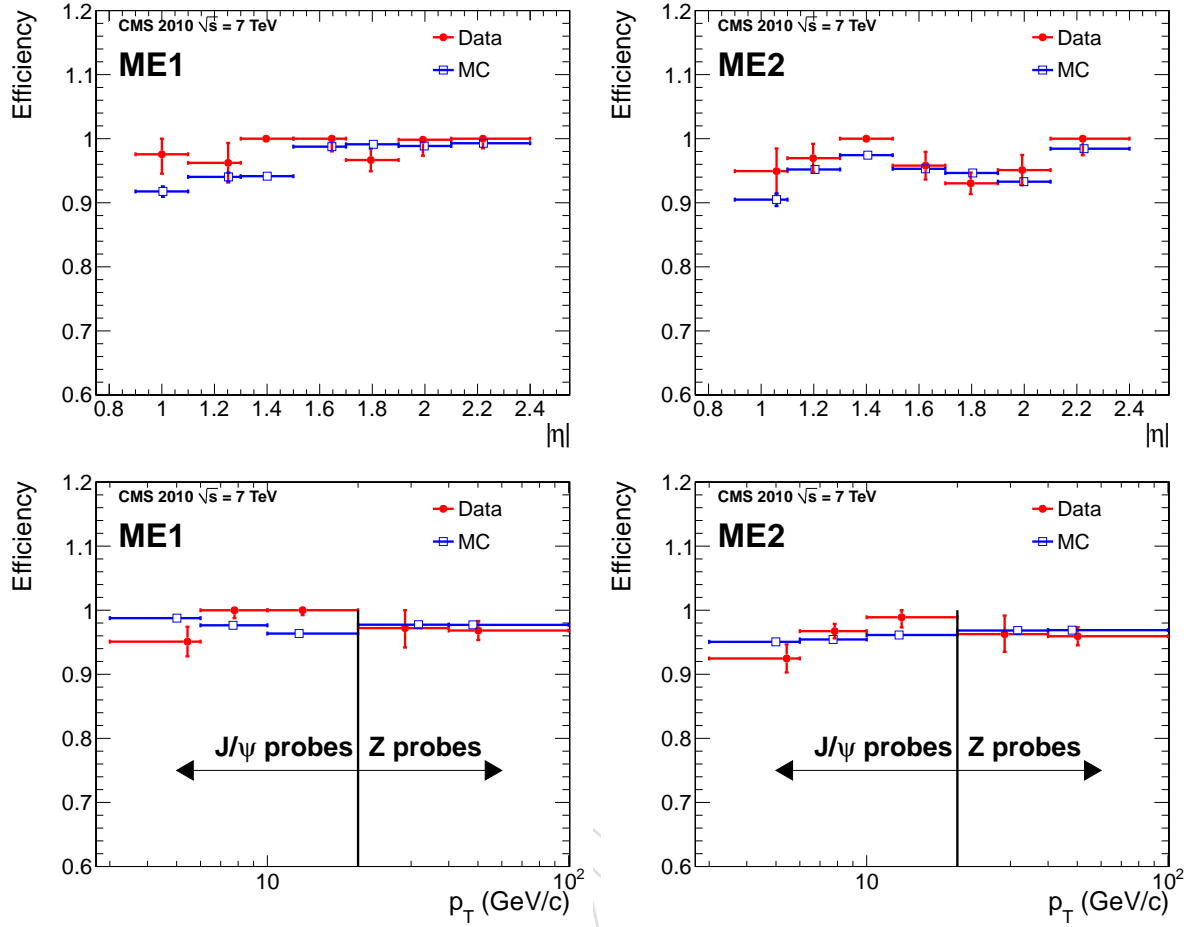


Figure 20: Comparison between measured (red) and simulated (blue) data of the segment reconstruction efficiency versus  $\eta$  (top), and  $p_T$  (bottom) for endcap CSC stations ME1 (left) and ME2 (right). The vertical line on the  $p_T$  distributions separates the ranges covered by probes originating from  $J/\psi$  and  $Z^0$  decays.

and local  $y$  is the coordinate orthogonal to it. The extrapolated track and the segment are considered to be matched if the distance between them in local  $x$  is less than 3 cm or if the value of the pull for local  $x$  is less than 4, where the pull is defined as the difference between the position of the matched segment and the position of the extrapolated track, divided by their combined uncertainties. The RMS width of residuals is shown in Fig. 21 as a function of the muon-station number, for the DT and CSC systems.

Tracker Muon reconstruction is more efficient than the Global Muon reconstruction at low momenta,  $p \lesssim 5$  GeV/c, because it requires only a single muon segment in the muon system, whereas Global Muon reconstruction is designed to have high efficiency for muons penetrating through more than one muon station and typically requires segments in at least two muon stations.

Thanks to the high tracker-track efficiency [33] and a very high efficiency of reconstructing segments in the muon system, about 99% of muons produced in pp collisions and having sufficiently high momentum are reconstructed either as a Global Muon or a Tracker Muon, and very often as both. Candidates found both by the Global Muon and the Tracker Muon approaches that share the same tracker track are merged into a single candidate. Muons reconstructed






Figure 21: RMS width of residuals of the local  $x$  position, given by the position of the muon segment with respect to the extrapolated tracker track, as a function of the muon station, for DT chambers in the barrel region (left) and CSCs in the endcap regions (right). Data are compared with MC expectations.

only as standalone-muon tracks have worse momentum resolution and less favourable collision muon to cosmic-ray muon ratio than the Global and Tracker Muons and are usually not used in physics analyses. Their resolution is at the level 10% in the barrel and 15-20% in the endcaps for muons with  $p_T$  up to 100 GeV/c.

For both Global and Tracker Muons with  $p_T < 200$  GeV/c, the momentum is determined by the tracker-only fit. The contribution from the muon stations becomes important only at high momenta. CMS has developed specialized algorithms for high- $p_T$  muon reconstruction and momentum assignment. As the muon passes through the steel of the magnet return yoke, multiple scattering and radiative processes can alter the muon trajectory. While the former is not so important for high-momentum muons, the latter can result in large energy losses and can also produce electromagnetic showers giving rise to additional hits in the muon chambers. As a consequence, the estimate of the muon momentum at the production vertex can be significantly different from its true value. Therefore, several different strategies for including information from the muon system have been developed and studied using cosmic rays [32]. These techniques either use hits only from the innermost station or retain only hits that, based on a  $\chi^2$  comparison, are compatible with the extrapolated trajectory. To further improve the resolution at high  $p_T$ , mainly by reducing the tails of the momentum resolution distribution, the *Tune P* algorithm chooses, on a muon-by-muon basis between these different fits.

Several techniques were developed to measure the muon momentum scale and resolution. For muons with  $p_T < 100$  GeV/c two different approaches, MuSclFit (Muon momentum Scale calibration Fit) and SIDRA (SImulation DRiven Analysis) described in MUO-10-004, making use of muons originating from J/ $\psi$  and Z decays. As can be seen in Fig. 22, the results obtained with the two methods agree within the uncertainties. The relative  $p_T$  resolution obtained using MuSclFit is found to be in the range from 1.3% to 2.0% for muons in the barrel and up to  $\approx 6\%$  for muons in the endcaps, in good agreement with the results obtained from simulation. The  $\sigma(p_T)/p_T$  averaged over  $\phi$  and  $\eta$  varies in  $p_T$  from  $(1.8 \pm 0.3(\text{stat.}))\%$  at  $p_T = 30$  GeV/c to  $(2.3 \pm 0.3(\text{stat.}))\%$  at  $p_T = 50$  GeV/c, again in good agreement with the expectations from simulation. Averaged over the whole acceptance, using the same techniques, the relative bias in the muon momentum scale is measured with a precision of better than 0.2% and is found to




Figure 22: Relative transverse momentum resolution  $\sigma(p_T)/p_T$  in data and simulation measured by applying the MuSclFit and SIDRA methods to muons produced in the decays of Z bosons and passing the Tight Muon selection. The thin line shows the result of MuSclFit on data, with the grey band representing the overall (statistical and systematic)  $1\sigma$  uncertainty of the measurement. The circles are the result of MuSclFit on simulation. The downward-pointing and upward-pointing triangles are the results from SIDRA obtained on data and simulation, respectively; the resolution in simulation was evaluated by comparing the reconstructed and “true”  $p_T$  once the reconstructed  $p_T$  was corrected for  $\phi$ -dependent biases. The uncertainties for SIDRA are statistical only and are smaller than the marker size.

be consistent with zero up to  $p_T$  values of 100 GeV/c. *should we comment about  $\phi$  and  $\eta$  dependent corrections needed for precision measurements*

High- $p_T$  muons are an important signature in many searches for new physics, so it is crucial that the performance of their reconstruction, which has some significant differences to that of lower- $p_T$  muons (such as an increased role of the muon system in momentum measurement and larger impact of showering) is well understood. Cosmic-ray muons that traverse the entire CMS detector can be used to evaluate the momentum resolution by comparing the momenta reconstructed independently in the upper and lower halves of the muon system, a procedure that was first applied to cosmic-ray muons collected in 2008 [32].

Figures 23(a) and (b) show the relative resolution as measured by the Gaussian width and the truncated sample RMS, respectively, for the tracker-only and global fits, and for the sigma-switch *[need to get rid of sigma switch]* and Tune P algorithms as a function of the  $p_T$  of the muon. The Gaussian width gives a measure of the core resolution, while the truncated sample RMS includes the effects of the tails of the distribution; both can be separately important for considerations of momentum resolution, possibly depending on the details of the physics analysis being considered.

*should add a paragraph on cosmic endpoint and quote 5% uncertainty per TeV in momentum scale.*

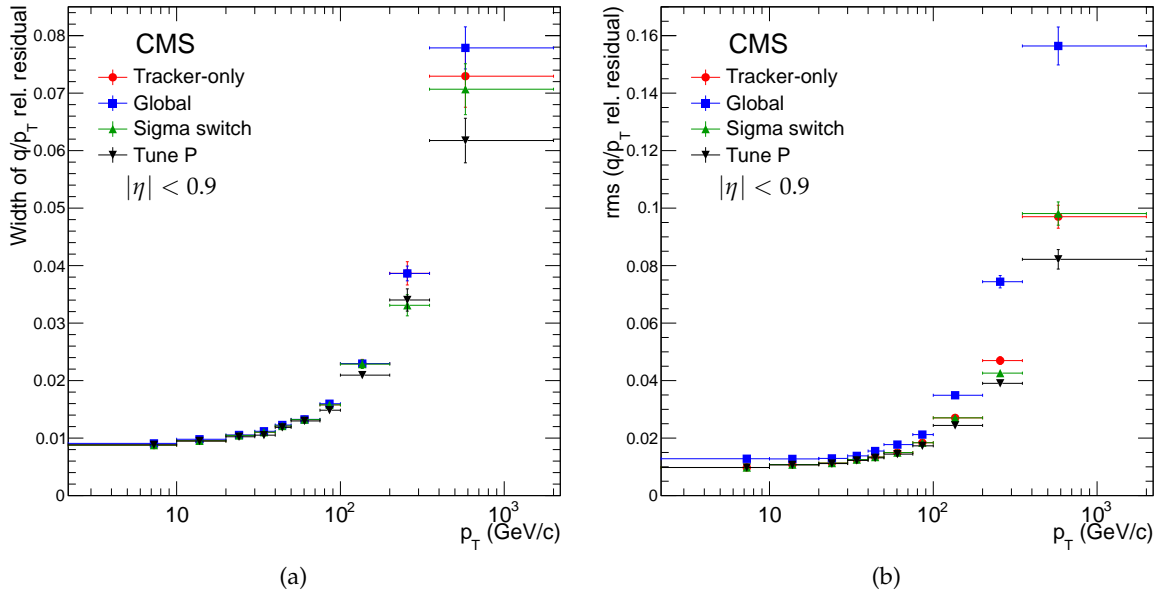
The combination of different algorithms provides a robust and efficient muon reconstruction. A given physics analysis can achieve the desired balance between identification efficiency and purity by applying a selection based on various muon identification variables. Three baseline muon identification algorithms, widely used in CMS analyses, are described here.

Table 10:  $\sigma$  position resolution for the RPCs with alignment.

Barrel		Endcaps			
Layer	$\sigma$ (cm)	Ring	$\sigma$ (cm)	Ring	$\sigma$ (cm)
RB1in	0.81	RE1/2/A	0.94	RE(2,3)/2/A	1.07
RB1out	0.90	RE1/2/B	0.88	RE(2,3)/2/B	0.96
RB2in	1.03	RE1/2/C	1.05	RE(2,3)/2/C	0.86
RB2out	0.99			RE(1,2,3)/3/A	1.11
RB3	1.06			RE(1,2,3)/3/B	1.28
RB4	1.32			RE(1,2,3)/3/C	1.10

Table 11: Local segment reconstruction efficiency for stations 1, 2, 3, and 4 of the barrel (DT) and endcap (CSC) muon systems .

	DT Efficiency (%)		CSC Efficiency (%)	
	Data	MC	Data	MC
Station 1	$99.2 \pm 0.4$	$98.05 \pm 0.03$	$98.9 \pm 0.9 \pm 0.8$	$97.8 \pm 0.1 \pm 1.1$
Station 2	$99.0 \pm 0.4$	$98.98 \pm 0.03$	$96.8 \pm 0.9 \pm 0.7$	$95.5 \pm 0.1 \pm 1.0$
Station 3	$99.1 \pm 0.4$	$99.08 \pm 0.04$	$96.8 \pm 0.9 \pm 1.1$	$94.1 \pm 0.1 \pm 1.7$
Station 4	$98.9 \pm 0.6$	$99.00 \pm 0.04$	$94.9 \pm 1.6 \pm 1.5$	$91.7 \pm 0.2 \pm 1.7$

Figure 23: (a) Widths of Gaussian fits of the distributions of the muon  $q/p_T$  relative residuals for the tracker-only and global fits, and for the output of the sigma-switch and Tune P algorithms, as a function of the  $p_T$  of the muon; (b) sample RMSs (truncated at  $\pm 1$ ) of the same distributions.

*this part needs to be adapted to 2012 situation*

- *Particle-Flow Muon selection.* The CMS particle-flow event reconstruction [34] combines the information from all subdetectors to identify and reconstruct individually particles produced in the collision. The resulting list of particles is then used to construct higher-level particle-based objects and quantities, such as jets and missing transverse energy. In order to identify Particle-Flow Muons, a selection is performed on all the muon candidates reconstructed with the standard algorithms described above. This selection has been optimized to identify muons in jets with high efficiency, keeping the misidentification rate from charged hadrons low. This is needed in order to avoid biases in jet and  $E_T^{\text{miss}}$  measurements coming from non-identified or misidentified muons. As a consequence, the Particle-Flow Muon selection has been designed to retain non-isolated muons, including the muons from hadron decays in flight, usually considered as a background in typical muon analyses. This is achieved by applying selection criteria, which differ in strictness depending on whether the muon candidate is isolated or not, and whether its momentum is compatible with the energy deposition in the calorimeters assigned to the candidate by the particle-flow event reconstruction. The details of the Particle-Flow Muon selection are described in Ref. [35].
- *Soft Muon selection.* This selection requires the candidate to be a Tracker Muon, with the additional requirement that a muon segment is matched in both  $x$  and  $y$  coordinates with the extrapolated tracker track, such that the pull for local  $x$  and  $y$  is less than 3. Segments that form a better match with a different tracker track are not considered. These additional requirements are optimized for low  $p_T$  ( $< 10 \text{ GeV}/c$ ) muons. This selection is used in quarkonia and B-physics analyses in CMS [36].
- *Tight Muon selection.* For this selection, the candidate must be reconstructed outside-in as a Global Muon with the normalized  $\chi^2$  of the global-muon track fit less than 10 and at least one muon chamber hit included in the global-muon track fit. In addition, its corresponding tracker track is required to be matched to muon segments in at least two muon stations (this implies that the muon is also reconstructed inside-out as a Tracker Muon), use more than 10 inner-tracker hits (including at least one pixel hit), and have a transverse impact parameter  $|d_{xy}| < 2 \text{ mm}$  with respect to the primary vertex. With this selection, the rate of muons from decays in flight is significantly reduced, at the price of a few percent loss in efficiency for prompt muons such as those from  $W$  and  $Z$  decays. The Tight Muon selection is used in many physics analyses in CMS, in particular in the measurements of inclusive  $W$  and  $Z$  cross sections [37, 38].

The Tight Muon transverse momentum and pseudorapidity distributions for  $p_T > 20 \text{ GeV}/c$  are compared to the expectations from the Monte Carlo simulation in Fig. 24. Muons from light-hadron decays are predicted to contribute less than 10%, while the hadron punch-through is suppressed to about 1%. The beauty contribution dominates up to muon transverse momentum of about  $30 \text{ GeV}/c$ , where the  $W$  contribution starts to prevail, leading to a shoulder in the falling  $p_T$  spectrum. The inclusive muon yield agrees with the expectations within a few percent up to a transverse momentum of  $50 \text{ GeV}/c$ . At higher momenta the leading processes are  $W$  and  $Z$  production, occasionally associated with hard jets. In this  $p_T$  region, the data agree with the predictions within 10%. This has been verified to be fully consistent with theoretical uncertainties related to missing higher-order QCD contributions, by comparing the MADGRAPH generator used to simulate  $W$  and  $Z$  with other Monte Carlo programs for the  $W(Z)$ +jets processes. In conclusion, given the known experimental and theoretical uncertain-

ties, the agreement between the data and simulation is satisfactory over the entire momentum range of  $p_T \lesssim 200 \text{ GeV}/c$ .

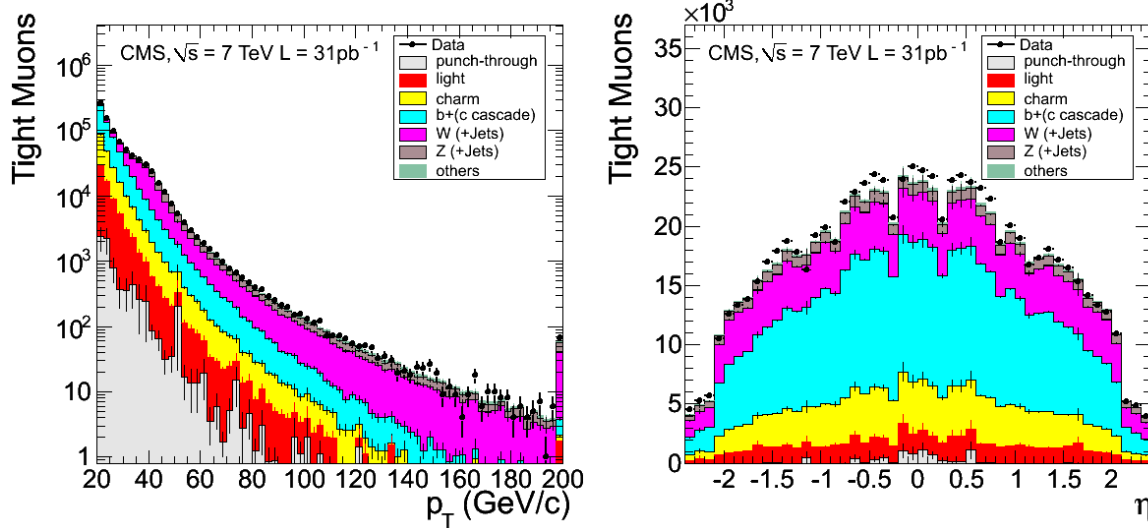


Figure 24: Distributions of transverse momentum (top left) and pseudorapidity (top right) for Tight Muons with  $p_T > 20 \text{ GeV}/c$ , comparing data (points with error bars) to Monte Carlo simulation broken down into its different components. The last bin in the  $p_T$  distribution includes the overflow.

We evaluate the efficiencies for prompt muons by applying a tag-and-probe technique to muons from  $J/\psi$  and  $Z$  decays. Using this technique it is possible to obtain almost unbiased estimates of the efficiencies of the different stages of muon trigger and offline reconstruction. Events are selected with strict selection requirements on one muon (the “tag” muon) and with a more relaxed selection on the other muon (the “probe” muon), such that the selection applied to the probe muon does not bias the efficiency that one wants to measure. The fraction of probe muons that passes the selection under study gives an estimate of its efficiency.

Figure 25 shows the muon efficiency  $\epsilon_{\text{rec+id}}$  given that a tracker track exists, measured using  $J/\psi \rightarrow \mu^+ \mu^-$  and  $Z \rightarrow \mu^+ \mu^-$  events. The results obtained from the data collected in the 2010 LHC data-taking period are compared with those from simulated events. The probabilities for particles other than muons to be identified as muons were studied using charged hadrons originating from decays of  $K_S^0$ ,  $\Lambda$  and  $\phi$ . These probabilities are 0.6% for Soft Muons, 0.4% for Particle Flow Muons and smaller than 0.2% for Tight Muons. *might want to update this part with 2012 data and selections*

In addition to the basic muon identification algorithms described above, additional selection criteria can be applied specific to the topology of the studied signal. An example is the transverse impact parameter cut which can help to separate prompt muons from muons originating from heavy-flavor decays. Another very useful quantity to increase the purity for prompt muons is the muon isolation. Several muon isolation algorithms were developed in CMS. Figure 26 shows the efficiency of the various isolation algorithms evaluated on muons with  $20 < p_T < 50 \text{ GeV}/c$  from  $Z$  decays as a function of the threshold on the corresponding isolation variable. Results obtained with the tag-and-probe (for all three isolation algorithms) and the LKT (for  $I_{\text{trk}}^{\text{rel}}$  and  $I_{\text{comb}}^{\text{rel}}$ ) methods are shown for both data and simulation. *maybe should drop LKT here.*

The good performance of the muon reconstruction and identification provides the necessary confidence in all elements of the chain from muon detection to muon analysis, which is essen-






Figure 25: Tag-and-probe results for the muon efficiency  $\epsilon_{\text{rec+id}}$  in data compared to simulation. Given that a tracker track exists, the plots show the efficiency as a function of muon  $p_T$  for Soft Muons (left), Particle-Flow Muons (middle), and Tight Muons (right) in the barrel and overlap regions (top), and in the endcaps (bottom). The measurement is made using  $J/\psi \rightarrow \mu^+\mu^-$  events for  $p_T < 20 \text{ GeV}/c$  and  $Z \rightarrow \mu^+\mu^-$  events for  $p_T > 20 \text{ GeV}/c$ . For  $p_T < 3 \text{ GeV}/c$ , only tracks with MIP signature are considered, in order to reduce the background.

1352 tial for searches for physics beyond the Standard Model as well as accurate Standard Model  
1353 measurements.

### 1354 3.4 Jet/MET

### 1355 3.5 Anomalous Signals in HCAL

1356 Several sources of anomalous signals in HCAL have been observed which can be misinter-  
1357 pretted as energy deposits. Electronics and detector noise observed primarily in the Hadronic  
1358 Barrel (HB) and Endcap (HE) occurs randomly and is independent of the beam conditions.  
1359 This type of noise occurring in the Hybride Photo Diode (HPD) and the Readout Box (RBX) is  
1360 collectively referred to as HPD/RBX noise. We also observe beam induced anomalous signals  
1361 recorded by the photomultipliers (PMTs) of the the Hadronic Forward (HF) calorimeter. The  
1362 source of this type of noise comes from particles interacting directly in the glass window, the  
1363 glass sides of the HF PMT, or the light collection fibers.

1364 Anomalous signals occuring later than the collision signal in HF were observed during op-  
1365 erations in 2010. The source of the signals was determined to come from scintillation light  
1366 produced in a sleeve used in the air core light guide of HF. In January 2011 the material was  
1367 replaced and this source of anomalous signal was removed.

1368 In HBHE the collision signal extends over several 25ns time samples. In 2011 the energy for




Figure 26: Isolation efficiency for muons from Z decays versus isolation efficiency for muons in the QCD-enhanced dataset described in the text, for tracker relative, tracker-plus-calorimeters relative, and particle-flow relative isolation algorithms. Muons are required to have  $p_T$  in the range between 20 and 50 GeV/c. The background rejection is limited by the 0.4% contamination from truly isolated muons.

HBHE was reconstructed using 4TS (100ns). With increasing luminosity and a bunch spacing of 50ns, it is possible for one channel to have energy deposits from two consecutive crossings. At the start of 2012, the integration window for the energy reconstruction was changed from using 4TS in HBHE to 2TS plus a signal containment correction in order to eliminate the contribution to the energy from energy deposits from the next collision crossing. For HF the energy reconstruction was changed from 2TS to 1TS in order to reduce the contribution from early hits from the next crossing. The HF signal is faster than for HBHE and is contained in one TS.

It is essential to remove anomalous signals so that they do not contribute to the measured energy attributed to collision event. Noise filtering algorithms have been developed to flag and remove anomalous signals from the reconstructed objects such as Jets and Missing Transverse Energy (MET). When many channels are identified as noise, the entire event can be flagged and removed during the user's analysis.

### 3.5.1 Noise Filters for HF

Particles interacting in the HF PMT can produce a signal arriving earlier than signals from Cherenkov light collected from the absorber. Such interactions also typically produce a signal in one PMT without a signal in the adjacent PMTs. The Cherenkov signal is narrow and is fully integrated within one 25ns time sample. The leading edge of the signal has been adjusted so that the collision signal is contained in the sample of interest and early PMT window events have part of the signal arriving in the earlier time sample. This allows us to use the signal sharing between the sample of interest and the previous sample to distinguish between anomalous and collision signals. In addition to anomalous signals arriving early we observe broad pulses that is attributed to showing in the region of the PMT.

The pulse shape discriminant  $PS = TS2/(TS1+TS2+TS3)$  is a measure of how well formed the pulse is and is used to flag reconstructed hits (rechits) as anomalous signals. Figure 27 shows the pulse shape variable as a function of the signal summed in the three time samples. The noise filters use this Additional requirement on the energy




Figure 27: Phases space spanned by the hit PS1 discriminator and 3TS energy value. The black contour gives the rough cut position. All hits found below the black contour are flagged as noise.

Particles that interact in the HF PMT produce a signal in either the long or short section of HF with no corresponding energy in the associated short or long section and no surrounding energy deposits. This property is exploited in the topological filter to flag anomalous hits.

Both the pulse shape and topological filter are used to flag rechits which are then removed from the reconstruction of calorimeter based objects. The HF noise filtering is applied by default during the offline reconstruction so that all physics analyses benefit from the cleaning.

### 3.5.2 Noise Filters for HBHE

One HPD consists of eighteen signal channel and one calibration channel. Four HPDs are grouped together in a Readout Box (RBX). Anomalous signals in HBHE has been observed that can effect anywhere from 1 to 72 channels. HBHE noise is classified by the number of channels within an RBX that have a signal above some threshold.

HPD/RBX noise is random and relatively stable within a factor of 2. For anomalous signals with  $E > 100\text{GeV}$  the rate of noise is about 3Hz.

The discriminant  $R45 = (TS4-TS5)/(TS4+TS5)$  is used to distinguish anomalous signals from in time collision signals in HB and HE. The rate of HBHE noise is low and R45 filter is used to flag the event as having noise and the user can then remove the event from their analysis. Figure 28 shows some example displays of events that were flagged by the HBHE R45 filter.

Filters that form a shape based discriminant were applied during 2012 data taking. The “fit-based” filters check whether the signal is consistent with a “spike” or a “flat” signal shape. Rechits having a pulse shape characteristic of noise are flagged and removed during the reconstruction of high level objects.




Figure 28: Event displays showing examples of HBHE noise flagged by the HBHE R45 filter. Left: Ion feedback. Middle: HPD Noise. Right: RBX Noise.




Figure 29: The MET distribution determined from the calorimeter before and after applying the HCAL noise filters.

### 3.5.3 HCAL Noise Filter Usage

HCAL noise filtering is applied at several stages. The HF topological filter which is less susceptible to changing LHC conditions is applied at the HLT. In order to further reduce the rate of noise triggered events it was necessary to apply a conservative version of the HBHE noise filter to the trigger paths that are most susceptible to HCAL related noise. More aggressive noise filtering is then applied during offline reconstruction allowing the possibility to tune the noise filters and reapply them during data reprocessing without any loss of data.

The HF pulse shape and topological filters are applied during offline reconstruction and used to flag and remove rechits from the reconstruction of calorimeter based objects such as Jets and Missing Transverse Energy. The fit-based filters are used to remove anomalous rechits in HBHE.

The impact on the MET distribution of the HCAL noise filters is shown in Figure 29.

Additional filters developed by the physics object group (POG) are also applied during the user's analysis to reject events. These filters remove the contribution from cosmic and beam halo Bremsstrahlung and some residual HBHE HCAL noise. Some exotic searches may not want to apply the standard noise filters.

K. Kousouris

## 3.6 Global event description

R. Cavanaugh

## 4 Level-1 Trigger, HLT

In the CMS design, the online event selection is achieved in two physical steps, the Level-1 (L1) Trigger and the High-Level Trigger (HLT). The L1 analyzes each crossing within a latency of  $\sim 3\mu\text{s}$ . It uses informations coming from the calorimeters and the muon subdetectors to reduce the event rate down to at most 100 kHz. The HLT operates on longer timescales and can make use of information from the tracking detectors. It accepts events at a rate of 300 – 350 Hz.

### 4.1 Level-1 Trigger

Editors: Luigi Guiducci, Jim Brooke

The Level-1 Trigger is a fully pipelined hardware data processor that uses information from the muon system and calorimeters to select interesting events for further consideration by the software based Higher Level Trigger. The L1 hardware uses FPGAs for fast, low latency, data processing, with interconnects based on fast serial copper and optical links. The system identifies trigger candidates (muons, EM candidates, jets and energy sums) that are used as input to a highly programmable menu of trigger algorithms, from which the final L1 Accept signal is generated and sent to the DAQ. As well as delivering the L1 Accept signal, the L1 hardware sends event data to the DAQ, for use in diagnostics, and seeding the Higher-Level Trigger.

Figure 30 shows the connectivity of subsystems that make up the L1 hardware. The front-end data from ECAL, HCAL, DT and CSC sub-detectors is used to generate trigger primitives. The calorimeters trigger primitives are ET sums over triggers towers with size  $0.0875\text{ eta} \times 0.0875\text{ phi}$ . For the DT and CSC sub-detectors, the trigger primitives consist of track stubs. The RPC system does not calculate trigger primitives, instead it performs track recognition directly using hit level information. The trigger primitives are sent to region processors (DTTF, CSCTF, RPCTF for muon systems and RCT for calorimeters). The regional processors feed global muon and calorimeter processors, which send output to the global trigger. The global trigger send a L1 accept pulse to the DAQ.

Figure 30: L1 Trigger hardware overview

*Highlight changes since TDR (and/or CRAFT paper?)*

#### 4.1.1 L1 Trigger Configuration during 2011

*Table with the main unprescaled triggers: single and double mu, single and double jet, single and double egamma, met, htt, ett, 2-3 cross triggers to be put here*

Beam pickups, called BPTX [ref], are connected to the Global Trigger input, and are used to implement gating of trigger algorithms, vetoing of non-colliding or colliding bunches and triggering on zero bias or single beam. In particular, with the LHC machine running at 50 ns minimum bunch spacing, the GT is configured in such a way to veto any physics trigger issued at the bunch crossing preceding a collision crossing, a technique called pre-BPTX- veto in jargon. The advantages of such technique are twofold: first, all out of time triggering inside LHC



bunch trains is suppressed, in particular the observed relatively large fraction of early triggering of the forward calorimeters at high luminosity. This results in a gain in trigger efficiency for the correct BX identification. Second, the veto allows a special configuration of the muon trigger to be implemented.

#### *L1 Muon Trigger Configuration*

A special configuration of the RPC track finder has been implemented to profit of the pre-BPTX-veto applied by the Global Trigger. Detector hits are extended to a duration of two bunch crossings and anticipated of one bunch crossing, enabling the RPC track finder system to be able to trigger both muons and muon-like slow particles at the correct collision bunch crossing. This technique extended significantly the sensitivity of searches for Heavy Stable Charged Particles.

RPC track finder firmware was also updated introducing new track patterns requiring hits from only 3 out of 6 RPC layers. The new patterns improve the trigger efficiency where detector gaps cause fewer hits to be associated to the muon trajectory, in particular around  $|\eta| 0.25$  (gaps between barrel wheel 0 and wheels  $\pm 1$ ).

CSC track finder  $p_T$  assignment was also updated, based on optimization from real data. This allowed a significant rate reduction while maintaining a very high efficiency for muons above the  $p_T$  threshold.

The transverse momentum measurement in muon chambers at  $2.1 < |\eta| < 2.4$  has poor resolution, due to the magnetic field geometry and detector coverage, resulting into up to 75% of the single muon trigger rate from this region. Single muon triggers with a restriction to  $|\eta| < 2.1$  were introduced, obtaining a large rate reduction for a relatively small loss of acceptance. The extra rate budget allowed a lower threshold (14 GeV) to be used at L1 and to relax the algorithm for  $p_T$  assignment in the Global Muon Trigger, resulting in an improvement of plateau efficiency by several percent.

Double muon triggers are not restricted in eta. The main double muon triggers apply either no  $p_T$  cut but tight quality cuts for triggering low  $p_T$  dimuons, for example L1\_DoubleMu0\_HQ used in quark-onia triggers, or a relatively high  $p_T$  cut on the leading muon only, for optimal efficiency with medium to high  $p_T$  muons, for example L1\_DoubleMu10\_open which is used for electroweak physics triggers. Quality cuts require the muon candidate to be identified with hits in more than two muon chambers and by different sub-detectors (DT/RPC and CSC/RPC), in an eta-dependent way, in order to improve the purity of the trigger.

Several other L1 single and double muon triggers of various thresholds are also present in the trigger menu, for covering different use cases in HLT seeding and for providing monitoring and commissioning triggers. As luminosity increased, several cross-triggers started to play a relevant role for many physics triggers, such as muon plus jets, muon plus e/gamma, etc.

#### *L1 Calo Trigger Configuration*

The calorimeter trigger delivers e/gamma and jet candidates, together with global energy sums, to the GT. E/gamma candidates are identified in the RCT using a 3x3 trigger tower sliding window algorithm, as described in [TDR]. E/gamma candidates are classified as isolated or non-isolated *details*. Jet candidates are identified in the GCT using a 3x3 trigger region sliding window algorithm. Jets are classified as central, tau or forward *details*. All candidates are sorted in the GCT and the four highest in each classification (non-iso EG, iso EG, central jet, tau jet, forward jet) are forwarded to the GT. Scalar and vector ET and HT sums are calculated in the GCT, where ET is summed over trigger regions, and HT is the sum over all identified jet candidates above an ET threshold, and sent to the GT. *Do we need to mention the HF rings? Have*

they been used?

Configuration and developments since startup. E/gamma corrections. Jet corrections. HF masking. Eta ranges. EG Isolation.

During the 2010/2011 data-taking, eta-dependent ET corrections were introduced to both E/gamma and jet candidates, to improve the response as a function of  $\eta$  and  $E_T$ . E/gamma corrections were derived initially from MC, then updated using corrections derived from data. The jet corrections were derived from MC. *Describe how these were calculated.* Tau isolation was added. *Details.* The energy sums are calculated over a programmable eta range, which was optimised to cover  $|\eta| < 3$ . *Details. EG Isolation*

#### 4.1.2 L1 Trigger Operation during 2011

The L1 trigger system is controlled and monitored by means of rack-mounted CPUs performing different tasks, namely i) interfacing to the electronics to configure programmable devices and read out status registers, ii) connecting to CMS databases [ref] to retrieve configuration data or to store monitoring information, and iii) displaying GUIs to trigger operators. Software applications are based on the Trigger Supervisor [ref] and XDAQ [ref] frameworks, enabling central control and monitoring through a semi- hierarchical communication network between the different control processes. A main Trigger Supervisor cell interfaces the CMS Run Control system, operated by the DAQ shifter, to the trigger subsystems, dispatching configuration requests and collecting status feedback.

Predefined trigger configuration modes (collisions, cosmics, circulating beams, special tests) prepared by trigger experts are available to the shift crew. Each mode maps to several database tables, containing the information about the trigger menu masks, prescale values, subsystem configurations.

The monitored quantities include results of actions (configure, pause, resume, etc), the status of software processes (exceptions, memory consumption, CPU time), the status of the hardware modules (information about the lock to the LHC clock, errors on trigger data links, scalers, etc). Hardware scalers are used to monitor the total L1 trigger rate, the rate of each algorithm or technical trigger, the amount of dead-time during data taking, and subsystem objects rates (e.g. rate of track segments in the muon detectors and of calorimeter trigger primitives). The information is timestamped and stored by subsystem applications to the CMS online database for offline analysis, while a central GUI, called L1- page, displays status information, alarms, error messages to the trigger operator, suggesting appropriate actions or expert contacts if needed. The instantaneous luminosity measurement provided by forward calorimeter scalers is available in the online applications, allowing the Global Trigger control application to suggest to the operator (or to apply automatically) the best set of prescale values to be used for the trigger configuration.

During the LHC run in 2011 the system operated very well, causing few disruptions of the CMS data taking. The trigger systems contributed to about 6% of the overall luminosity lost by CMS due to problems while LHC was delivering proton-proton collisions (stable beams). There were 104 events of interruption of the data taking caused by one of the L1 Trigger systems, including both software and hardware problems. These events accounted for the loss of about 7 hours to be compared to a total of about 1371 hours of stable beams from the LHC. In terms of luminosity, about  $20 \text{ pb}^{-1}$  were lost due to operation problems.

### 4.1.3 L1 Trigger Monitoring and Data Certification

Several quantities are monitored in the CMS Data Quality Monitoring framework [ref], looking at physics events from the data acquisition stream. Typically, each L1 trigger subsystem (DTTF, CSCTF, RPCTF, RCT, GCT, GMT, GT) sends trigger event data, consisting in a snapshot of module's input and output information, to the DAQ path. A large set of histograms and summaries are available to the trigger operator, who can check in quasi-real-time the quality of the trigger data. The main observables used to qualify the L1 trigger during the data taking are described hereafter.

#### *Trigger cross sections*

Instantaneous luminosity as measured by the forward calorimeter (HF) and rates of all L1 trigger bits are available in the data stream. The cross sections of several trigger objects are calculated for each luminosity section (interval of data taking corresponding to about 23.31 seconds) and compared to expected values retrieved from CMS databases as a parametrization of the instantaneous luminosity. Thus the behaviour of each trigger object with respect to luminosity is certified with high time granularity.

#### *Emulator-data comparison*

Bit-wise emulators are implemented in CMSSW, [ref] the CMS analysis framework. The functioning of hardware modules is thus exactly checked for events available to the DQM processing. Each module's output data is compared with the output of the emulator, fed with the module's input data. The expected agreement is typically 100%; exceptions accounting for some level of discrepancy, due to known hardware or emulator problems, were allowed, by tuning a threshold for considering the emulator test successful.

#### *Synchronization*

Each trigger system is expected to be synchronized to the actual LHC collisions. The synchronization was monitored online comparing the BX information of a given object with respect to the beam crossing BX, identified by beam pickups (BPTX). Both early and late trigger objects can be observed in an unbiased way thanks to the presence of the pre-beam-crossing veto. Due to the low overlap between different objects and to the low statistics available online, this test is performed with a lower time-granularity, adding the information from consecutive luminosity sections until a few percent precision is reached.

#### *Occupancy*

The occupancy of trigger objects in the (eta, phi) coordinates is tested for hot or cold/dead regions, using the symmetry features of the detectors. Failing tests are flagged and problematic regions and objects are identified automatically. As the synchronization test, the occupancy test is performed adding information from several consecutive luminosity sections to achieve the desired precision.

All quality tests and other histograms are used by the trigger operator and experts to assert results for the CMS Data Certification process. Problems could be identified and data taking periods excluded from the physics analysis samples, down to a granularity of the single luminosity section. In the full 2011 run, 29  $\text{pb}^{-1}$  were declared bad data by the L1 Trigger certification procedure, approximately 0.6% of the luminosity recorded by the CMS experiment. A single problem on beam pickups accounted for a loss of 19  $\text{pb}^{-1}$ , while other sporadic problems affected typically the configuration of electronics or the data links between different processing modules.

#### 4.1.4 L1 Muon Trigger Performance

*to be written. Results from Laria. Efficiency analysis with tag and probe*

The muon trigger efficiency as a function of  $p_T$  is shown in Figure 31.

Figure 31: Muon trigger turn on curves for a range of thresholds

#### 4.1.5 L1 Calorimeter Trigger Performance

The e/gamma, jet and energy sum trigger efficiency is shown as a function of  $E_T$  in Figures 32,33,and 34.

Figure 32: E/gamma trigger turn on curves for a range of thresholds

Figure 33: Jet trigger turn on curves for a range of thresholds

## 4.2 High Level Trigger

editors: Emmanuelle Perez, Christos Leonidopoulos

The HLT hardware consists of a single processor farm, the “Event Filter Farm”, (EvF) described in section 5. In contrast to traditional three-level trigger systems, where the HLT selection starts with a “Level 2” step based on the detector information around “regions of interest” indicated by the L1, the CMS Filter Farm can use full granularity data from the whole detector, and the selection can be based on sophisticated offline-quality reconstruction algorithms. Due to the high input rate of  $\sim 100$  kHz that the filter farm needs to sustain, this requires significant CPU resources and the algorithms that run during the HLT processing must be optimized for performance in order to minimize dead-time. With the 2011 configuration of the EvF, the CPU power available allowed L1 input rates of 100 kHz to be sustained for an average HLT processing time of up to  $\sim 90$  msec per event. A short description of the algorithms used in the EvF and of their performance is given in the next paragraphs. Each HLT trigger is implemented as a sequence of reconstruction and selection steps of increasing complexity, reconstruction refinement and physics sophistication: selections relying on information from the calorimeters and the muon detectors reduce the rate before the CPU-expensive tracking reconstruction is performed. The reconstruction modules and selection filters of the HLT use the software framework that is also used for offline reconstruction and analyses.

### 4.2.1 Triggers based on jets and global event quantities

**4.2.1.1 Jet triggers** At the HLT, calorimeter jets are reconstructed using the “antikt” algorithm with a distance parameter of 0.5. The inputs to the jet algorithm are calorimeter towers, which are constructed from the energy deposited in projected HCAL cells and corresponding projected ECAL crystals that satisfy certain threshold requirements. Loose jet identification cuts are applied that increase the robustness of the jet candidate identification with respect to calorimeter noise. Jet energy corrections are applied to equalize the jet response with respect to the jet pseudo-rapidity.

A series of prescaled jet triggers allows the inclusive jet spectrum to be measured down to low  $E_T$  and provides control samples for efficiency studies. High threshold inclusive jet triggers are used in searches for new phenomena. High multiplicity jet triggers can run unprescaled with a relatively low threshold. For example, a trigger requiring four (six) jets with  $E_T$  above 80 GeV (45 GeV) has a rate of 6 Hz (2 Hz) at a luminosity of  $5 \cdot 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ . These triggers are also

Figure 34: Energy sum trigger turn on curves for a range of thresholds

used in searches for new physics. Several triggers that combine a jet requirement with a lepton requirement have also been designed and are used in several analyses.

**4.2.1.2 Particle-flow reconstruction at the HLT** Starting in 2011, CMS has started to exploit the “particle-flow” (PF) based reconstruction (see Sec. 3.6) in the Online Selection offering improved resolution and lower trigger rates without loss of efficiency. To comply with the CPU constraints of the HLT, the PF algorithms, in particular the steps that perform the “iterative tracking”, had to be simplified with respect to their offline counterparts. A pre-filtering of events based on calorimeter based quantities is usually made before running the PF reconstruction, in order to limit the rate at which the latter is executed down to a few kHz. Once PF particles are reconstructed, they can be used as input to the jet algorithm to build “PF jets”. Figure 35 (left) shows example efficiencies of single jet triggers using online PF jets. It can be seen that the turn-on of these efficiency curves is very sharp. The usage of PF reconstruction in the HLT has continuously increased since it was introduced in 2011. Most of the triggers are now making use of it to reconstruct jets and the global event quantities described below.

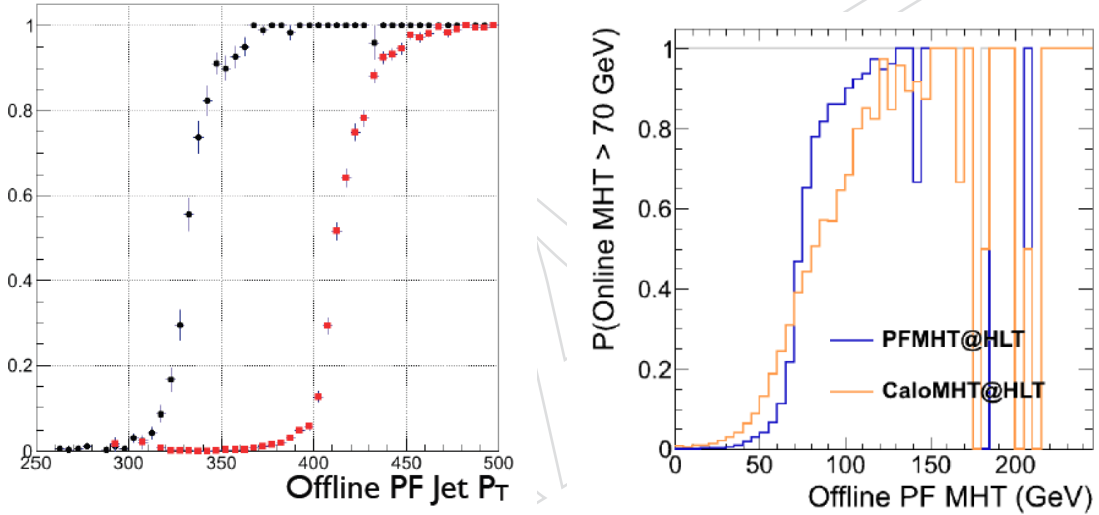


Figure 35: Left: trigger efficiency of an inclusive trigger requiring an online PF jet with  $E_T$  above 320 GeV or 400 GeV. Right: Trigger turn-on curves of a trigger requiring  $MHT > 70$  GeV, using the standard calorimetric reconstruction (orange curve) and using the Particle-flow reconstruction (blue curve).

**4.2.1.3 Triggers based on global event quantities** The missing transverse energy and the total transverse energy can also be used in the online event selection. The missing transverse energy,  $E_T^{\text{miss}}$ , can be reconstructed from the calorimeter towers (“MET”), from the vectorial sum of the transverse energies of calorimeter jets above a given threshold (“MHT”), or from all particles found in the particle-flow reconstruction (“PFMET”). The latter has a much better resolution and includes deposits from muons, which are not accounted for in a calorimeter-based approach. For example, a trigger requiring the calorimetric  $E_T^{\text{miss}}$  to be above 200 GeV reaches an efficiency of 95% when the offline  $E_T^{\text{miss}}$  is larger than  $\sim 270$  GeV. In contrast, a trigger requiring PFMET to be larger than 180 GeV is 95% efficient already for offline  $E_T^{\text{miss}}$  values of  $\sim 200$  GeV. Figure 35 (right) shows the efficiency of a trigger requiring a missing transverse



energy above 70 GeV, calculated from calorimeter jets or from the PF particles; the turn-on is much sharper in the latter case.

The total transverse energy,  $H_T$ , is calculated either from the calorimeter jets which have an  $E_T$  above a given threshold (“HT”), or from the scalar sum of the transverse energies of the PF jets above a given threshold (“PFHT”). Due to the better resolution provided by the PF reconstruction, lower thresholds can be sustained for triggers using the PF-based calculation. At a luminosity of  $5 \cdot 10^{33} \text{cm}^{-2}\text{s}^{-1}$ , a trigger requiring that PFHT be larger than 650 GeV runs unrescaled with a rate of  $\sim 10$  Hz.

Other triggers can combine a requirement on (PF)HT and on (PF)MET, or make use of analysis-oriented variables, such as the “ $\alpha_T$ ” variable which allows a good separation of events with genuine  $E_T^{\text{miss}}$  from those where a mismeasurement fakes a large  $E_T^{\text{miss}}$  value, or the “razor” variable which best exploits the kinematic properties of specific processes.

These triggers are extensively used in searches for supersymmetry and other new phenomena.

**4.2.1.4 Pile-up mitigation** Triggers that rely solely on jets or on HT are very sensitive to the number of overlaid proton-proton interactions (pile-up). The left plot in Fig. 36 shows the rate of a trigger requiring HT above 350 GeV and MHT above 110 GeV as a function of the instantaneous luminosity. The measurements at  $\mathcal{L} = 2 - 3 \cdot 10^{33} \text{cm}^{-2}\text{s}^{-1}$  were made during standard collision runs taken in summer 2011, where the average pile-up was 10 – 15, while the rates in the rightmost part of the plot were measured during a special high intensity run, taken in October 2011 with 10 colliding bunches. This run had an average pile-up of  $\sim 31$  and would correspond to  $\mathcal{L} \sim 5.5 - 7 \cdot 10^{33} \text{cm}^{-2}\text{s}^{-1}$  with 1330 bunches. These latter measurements markedly deviate from a linear extrapolation of the rates observed at lower luminosity, indicated by the blue line, and show the typical dependence of purely hadronic triggers upon pile-up. To mitigate this pile-up effect, an algorithmic subtraction of the “pile-up noise” from the jet transverse energies has been deployed. The algorithm is similar to the “FastJet” subtraction described in [? ], but its implementation had to be simplified in order to comply with the timing constraints of the HLT. The  $E_T$  density due to pile-up is first estimated, on an event-by-event basis, from all the jets reconstructed in the event; it is then used to correct the transverse energy of the jets. Figure 36 shows that the rate of the previously considered trigger is dramatically reduced when HT is calculated from pile-up corrected jets. Although some deviation from linearity still remains, the pile-up correction brings the rates down to a manageable level. In 2012, this correction is used in all triggers employing calorimetric jets<sup>2</sup>.

**4.2.1.5 b-tagged jets** In 2011, lifetime-based  $b$ -tagging at the HLT used the impact parameters (IP) of tracks associated to a jet. The primary vertex is first determined in three dimensions, using tracks reconstructed with the pixel detector, with a resolution of typically 100  $\mu\text{m}$  in both the longitudinal and the transverse direction. A jet is  $b$ -tagged if two tracks associated to the jet have a large impact parameter, with a significance above a given cut. This requirement is first applied on the pixel tracks associated to the jets. When it is satisfied, tracks are reconstructed using the pixel and the strip detectors, and a similar requirement is made on the full tracks associated to the jet. To ensure a good trigger efficiency with respect to jets that are  $b$ -tagged offline, it is important that the algorithm select the proper primary vertex (this is done by maximizing the  $\sum P_T^2$  of pixel tracks associated to it). In the presence of high pile-up, the probability that the selected vertex comes from a PU interaction is not negligible. In 2011,

<sup>2</sup>Triggers based on PF jets or PFHT are much more robust against pile-up. However, as mentioned above, a filtering based on calorimeter quantities is first needed since the PF reconstruction can not be run online at too large a rate. The triggers based on PF objects use the pile-up corrected calorimeter jets in this first step.




Figure 36: Reduction of the HLT rates via the FastJet subtraction.

the mis-characterization of a PU interaction as primary vertex limited the plateau of the online  $b$ -tagging efficiency to about 90%.

Another limitation of this algorithm comes from the fact that the full reconstruction of pixel tracks is time consuming, and uses a large fraction of the HLT CPU resources unless the rate at which it is run is first reduced, by e.g. requiring the presence of high  $E_T$  jets. This has been overcome in 2012 via the deployment of a new algorithm. A fast, coarse reconstruction of the primary vertex is first performed by histogramming the  $z$ -position of pixel clusters that are geometrically matched, in azimuth, with a calorimeter jet. The subsequent reconstruction of pixel tracks is limited to those that are constrained to this primary vertex, which considerably reduces the combinatorics. The vertex is then redetermined using these pixel tracks, and the rest of the IP-based  $b$ -tagging algorithm follows as described above. The reduced CPU requirements for the determination of the primary vertex allows for secondary vertices to be searched. In 2012, triggers filtering on the presence of a secondary vertex are used, in addition to the original triggers employing the IP-based approach.

#### 4.2.2 Electron and photon triggers

A regional reconstruction of the energy deposited in the crystals of the ECAL is performed, around the L1 EM candidates, followed by the building of “superclusters” (SC) using the same algorithms as used in the offline reconstruction. Identification criteria are applied to the superclusters that pass a given  $E_T$  threshold, based on a cluster shape variable (the RMS of the width of the shower) and on the ratio  $H/E$  of the energy deposited in the HCAL within a cone around the SC to the SC energy. Moreover, isolation requirements can be made with respect to the additional energy deposited in the ECAL and the HCAL, in a cone of  $R \sim 0.3$  around the EM candidate. The HLT uses several “working points” for these calorimeter-based requirements. They reduce the rate by a factor of typically 3 – 4, reaching  $\sim 10$  for the tightest working point used in 2011 – 2012. The thresholds are such that, after this set of calorimetric criteria, the rates of electron candidates are  $\mathcal{O}(1)$  kHz.

The next step of the online selection of electrons involve the tracker. It starts with a “pixel-matching step”, which uses the energy and position of the super-cluster to propagate back through the magnetic field (under both charge hypotheses) and to search for compatible hits in the pixel detector. Full tracks are then reconstructed from the resulting pixel seeds. Tim-

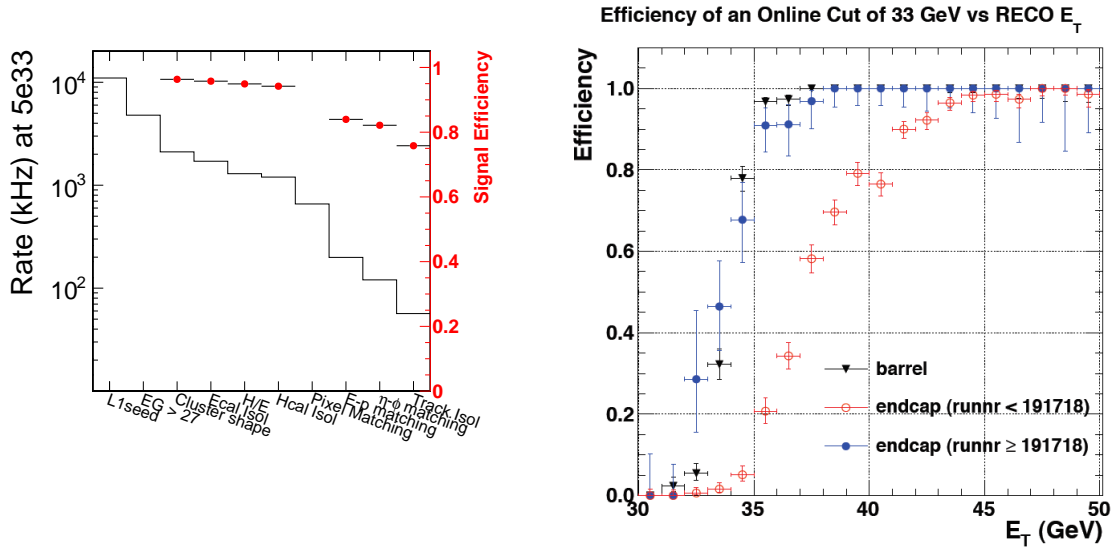


Figure 37: Left: Performance of the various steps of an example electron trigger. The rate is shown as the black histogram (left scale); the red symbols show the efficiency on genuine electrons (right scale). Right: Efficiency of an online  $E_T$  cut of 33 GeV as a function of offline electron  $E_T$ , in the barrel and in the endcap, before and after the deployment of online laser corrections.

ing constraints prevent the usage of the offline tracking algorithms and a simple Kalman filter technique is used<sup>3</sup>. A track should be reconstructed with a momentum compatible with the SC energy; its direction at the last tracker layer should match the SC position in  $\eta - \phi$ . These selection criteria further reduce the rates by another factor of  $\sim 10$ . Finally, an isolation requirement with respect to the remaining tracks reconstructed around the electron candidate can be applied.

With tight requirements, an inclusive electron trigger with a threshold of 27 GeV can be kept unprescaled at a luminosity of  $5 \cdot 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ , with a rate of  $\sim 50 \text{ Hz}$ . Figure 37 (left) shows how the rate is gradually reduced by the various filtering steps of this trigger, along with the efficiency on genuine electrons. Triggers with lower thresholds and relaxed identification requirements run prescaled and collect statistics to be used as control samples. Other triggers that require another object in addition to the electron candidate operate at higher efficiency working points for the electron selection. For example, double electron triggers are of crucial importance to many analyses. A suite of dielectron triggers with thresholds of 17 GeV and 8 GeV has been designed to collect signal events and control samples, with loose isolation criteria on both legs. These would allow to trigger on  $H \rightarrow WW \rightarrow e\bar{\nu}\nu$  events with an efficiency larger than 98% for a Higgs mass of 130 GeV, for a rate of less than 10 Hz at a luminosity of  $5 \cdot 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ .

The HLT selects photon candidates with the calorimetric requirements described above. Isolation criteria, based on calorimeter information and on the tracks reconstructed in a cone around the photon candidate, can also be applied. Inclusive photon triggers of low- $E_T$  are used for monitoring and QCD studies, while high- $E_T$  single photon triggers are used in searches for new physics. Di-photon triggers, used to search for a Higgs boson decaying into a pair of photons, are of particular interest. For these, another variable is used to complement the isolation

<sup>3</sup>Since 2012, it is complemented by the Gaussian-Sum Filtering algorithm.

criteria:  $R_9 = E_{3 \times 3} / E_{SC}$ , where  $E_{3 \times 3}$  denotes the energy deposited in a small window of  $3 \times 3$  crystals around the most energetic crystal in the SC. Requiring that  $R_9$  be larger than 0.8 (nominal cut) or 0.9 (tight cut) efficiently selects unconverted photons even in the presence of large pile-up. The standard isolation variable described above is less robust against pile-up, since it sums up the energy deposited in a larger cone area<sup>4</sup>. The current strategy of the diphoton triggers is to recover high  $R_9$  photons that fail the standard isolation criteria. Thresholds of 26 GeV and 18 GeV are applied on the transverse energies of the two photons. An additional cut on the transverse mass of the  $\gamma\gamma$  pair of 60 GeV allows this trigger to run unprescaled at a luminosity of  $5 \cdot 10^{33} \text{cm}^{-2} \text{s}^{-1}$ , with a rate of about 20 Hz. The trigger efficiency for the selected events of the  $H \rightarrow \gamma\gamma$  analysis based on the 2011 data was larger than 99%.

An important improvement that has been deployed in the trigger in 2012 is the use of corrections for radiation induced changes in the transparency of the crystals in the Endcap ECAL. These changes are determined via the laser monitoring system of the ECAL. A new set of corrections is deployed every week. Figure 37 (right) shows that the introduction of these corrections in the trigger has considerably improved the performance of the electron trigger in the endcap, which is now similar to the performance achieved in the barrel.

### 4.2.3 Muon triggers

The muon triggers at CMS combine information from the muon and the tracker subdetectors to identify muon candidates and determine their transverse momenta,  $p_T$ . A Level-1 muon candidate is first used as a seed to reconstruct a standalone muon track in the muon system. A full track fit is performed with a  $p_T$  measurement relying exclusively on the muon system which is used for the first filtering stage at the HLT. The next step of the reconstruction uses seeds in the silicon tracker that are generated in the region around the standalone muon track, and tracks are reconstructed in the tracker. If the tracker track and the standalone muon track are compatible with each other, a global fit combining tracker and muon hits is performed, yielding a muon candidate with an improved  $p_T$  measurement on which the final selection can be made.

Figure 38 shows the turn-on efficiency curves for a single muon trigger for the barrel ( $|\eta| < 0.9$ , left) and the endcap ( $0.9 < |\eta| < 2.1$ , right) regions. The single muon trigger reaches efficiencies above 95% (90%) in the central (forward) detector. Several muon triggers have been deployed and used for a wide range of physics analyses. As an example, a single muon trigger with  $p_T > 40$  (24) GeV confined in the  $|\eta| < 2.1$  region without (with) an isolation requirement gives a rate of 15 (28) Hz at a luminosity of  $5 \cdot 10^{33} \text{cm}^{-2} \text{s}^{-1}$ . A double muon trigger in the  $|\eta| < 2.4$  region with asymmetric (17, 8) GeV thresholds on the two legs has a rate of 7 Hz at the same instantaneous luminosity. In addition to the above “baseline” single and double muon triggers, CMS has designed specialized dimuon triggers aiming at capturing a significant number of onia and  $B$  physics topologies. This is typically achieved by dropping the requirement of a good track segment in the muon system for the second muon, thus allowing to reduce the effective kinematic threshold on the asymmetric triggers. Additional filtering for reducing the low- $p_T$  background rate can include mass cuts on the dimuon system, requirements on the angle between the two muon candidates, etc. Fig. 39 demonstrates the wide range of the dimuon mass spectrum that CMS captures with its muon triggers, including low-mass ( $\omega$ ,  $\phi$ ), onia ( $J/\psi$ ,  $\psi'$ ,  $B_S$ ,  $Y$ ) and electroweak ( $Z$ ) or beyond ( $Z'$ ) resonances.

<sup>4</sup>Applying a “FastJet” correction to the isolation variables to subtract the energy that is due to the underlying interactions is under study.

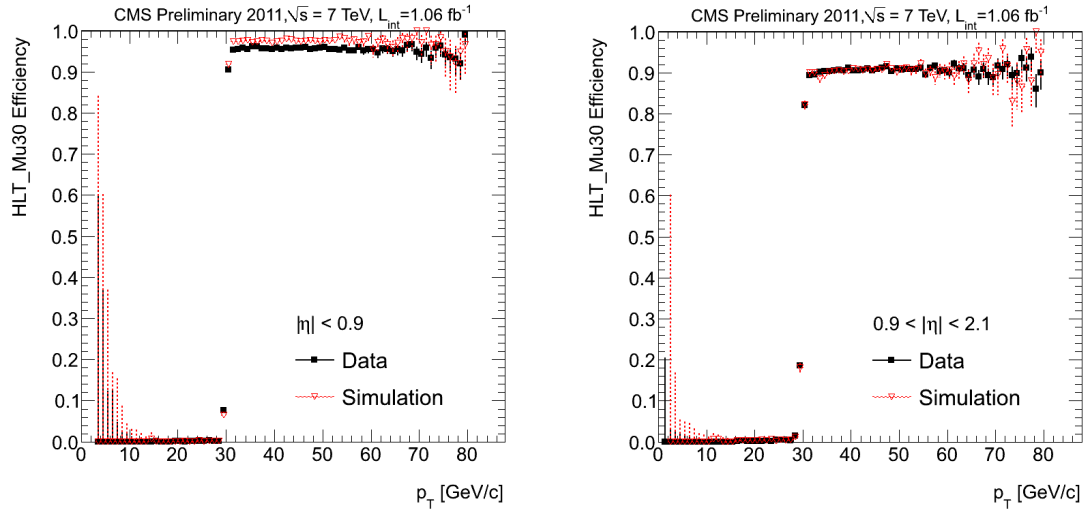


Figure 38: Efficiency curves for a single-muon trigger with  $p_T > 30$  GeV as a function of  $p_T$  of offline, good quality muons for the barrel ( $|\eta| < 0.9$ , left) and the endcap ( $0.9 < |\eta| < 2.1$ , right) regions. The single muon trigger reaches efficiencies above 95% (90%) in the central (forward) detector. NB: reproduce plots without the simulation curves?




Figure 39: The versatility of the CMS trigger allows for a wide range of the dimuon mass spectrum to be recorded using single and double muon triggers.



#### 4.2.4 Tau triggers

Tau triggers are used for a large number of Higgs and other Exotic searches. CMS has designed a suite of single and double tau triggers, as well as combinations of tau objects with  $E_T^{\text{miss}}$ , missing  $H_T$ , jets or leptons. Level-1 tau or central jet candidates are used as seeds. A calorimetric jet with  $R = 0.5$  is used as the first filtering stage at the HLT to reduce the background rate before a second filtering is applied requiring isolation calculated with pixel tracks. The subsequent filtering steps involve more sophisticated PF algorithmic techniques, including clustering of energy deposits in the ECAL and HCAL, and PF-based tau identification employing PF jets, leading-track finding and charged-track and photon isolation. For the tau-plus-lepton triggers the extra lepton requirement allows the simplification of the above sequence and the first filtering based on calorimetric jets is dropped.

A design decision to use track-based isolation for the tau triggers has resulted in an increased signal efficiency and only mild dependence upon pileup, as can be seen in Figure 40, left. Figure 40, right shows the turn-on efficiency curve for a single tau trigger with loose isolation and  $p_T > 20$  GeV. The trigger reaches efficiencies above 92%.




Figure 40: Efficiency of single-tau trigger with loose isolation and  $p_T > 20$  GeV as a function of the number of reconstructed primary vertices (left) and as a function of the  $p_T$  of offline, good-quality reconstructed tau objects (right).

Example rates for some of the tau triggers deployed by CMS are 2 Hz for a  $\tau + E_T^{\text{miss}}$  trigger with  $p_T(\tau) > 35$  GeV and  $E_T^{\text{miss}} > 70$  GeV, 14 Hz for a double-tau trigger with  $p_T > 30$  GeV and medium isolation, and 8 Hz for a  $\tau + \mu$  trigger with  $p_T(\tau) > 20$  GeV,  $p_T(\mu) > 17$  GeV and  $|\eta_\mu| < 2.1$  at a luminosity of  $5 \cdot 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ .

#### 4.2.5 Calibration triggers

Calibration techniques typically require very large statistics, however the full detector information is generally not needed. In order to record the appropriate datasets for calibration without saturating the data-taking bandwidth, dedicated triggers have been designed that record events with reduced content, such that the typical event size is around 1 – 5 kB containing only the detector information that is relevant for the calibration algorithms. These events are written to a special data stream.

For example, two such triggers select events that are used for the calibration of the ECAL. The first one collects minimum bias events and only the ECAL “hits” are recorded. By exploiting the  $\phi$  invariance of the energy deposition in physics events, this sample allows the intercalibration within a  $\phi$  ring of the electromagnetic calorimeter to be performed. The second ECAL trigger reconstructs  $\pi^0$  and  $\eta$  meson candidates decaying into two photons and only the ECAL hits associated to these photons are kept.

## 5 DAQ

F. Meijers

## 6 Software

### 6.1 High Level trigger

#### 6.1.1 Streams and Primary DataSets

### 6.2 Data Handling

#### 6.2.1 Tier0 workflows

#### 6.2.2 Tie1 workflows

### 6.3 Reconstruction software

#### 6.3.1 Introduction

The effort for the Reconstruction software in CMSSW started in early 2005, soon after the decision to move to the new Event Data Model. While in the first version most of the reconstruction algorithms were simply adapted from the previous CMS recostruction code (ORCA[? ]), quite soon new development were inserted. The transitioning phase, in which CMSSW was under development and not usable for physics studies, lasted up to Summer 2006, when the CMSSW version 1.X was successfully used to for the Computing, Software and Analysis challenge 2006 (CSA06[? ]). The Reconstruction software used for early LHC data (the 900 GeV run, late 2009), is quite far from this initial code: most of the algorithms have been either changed or deeply optimized in the meantime, and performance of time-critical pieces of code (tracking, for example) have been rehauled. Current (2012) Reconstruction code performance exceed the estimates present in the Computing TDR ([? ]), even with 2012 running conditions harsher than expected.

### 6.4 Algorithms improvement for 2012 data taking

LHC luminosity increased during 2011 from  $3^{32}\text{cm}^{-2}\text{s}^{-1}$  to more than  $3^{33}\text{cm}^{-2}\text{s}^{-1}$ . This directly translated in a number of Minimum Bias interactions overimposed to the signal event varying from less than 5 to more than 20. After the summer 2011 it started to become clear that during 2012 the number of superimposed interactions would have nearly double in 2012, with an LHC luminosity exceeding  $7^{33}\text{cm}^{-2}\text{s}^{-1}$ . The lessons learnt during 2011 made soon very clear that the CMSSW release used at the end of 2011, CMSSW\_4.4.X, would have not survived in such an environment, without a deep optimization of the most time consuming and memory hungry algorithms. A task force was launched to explore the changes needed to accomplish a CMSSW version able to cope with 2012 conditions, assessing both the technical changes, and the (eventual) price to pay in terms of physics performance. In the last period of 2011 data taking, a special LHC fill was put into collision, with fewer than usual bunches in the

machine, but with a number of protons per bunch similar to that expected in 2012 (The so-called “high PileUp run”). This allowed to perform performance tests on datasets virtually identical to those expected during 2012.

#### 6.4.1 Algorithmic Improvements

Changes in the algorithms used to reconstruct data can easily lead to a decrease in resources used during the processing. The easiest way to obtain less resource hungry algorithms is relatively easy: decrease physics performance (for example, perform tracking reconstruction increasing the  $p_T$  threshold). The task force decided instead to try and follow a different direction: only changes which do not impact physics performance were accepted. In this way, the astounding performance of the CMS detector, as shown in years 2010 and 2011, is completely maintained for the 2012 run.

The most time consuming algorithms in CMSSW\_4.4.X are those related with CMS tracker reconstruction (seeds, tracks, vertices, ...), and were natural candidates for the optimization. Track reconstruction was attacked both at the input stage (tracker hits were redesigned to use less memory, and with a shorter virtual hierarchy), initial stage (seeding, using cluster shapes to discard early combinatorial fakes), during trajectory building (by optimizing the settings for the 7 steps iterative tracking), and for special cases taking previously a lot of time to reconstruct (special algorithms to reconstruct low  $p_T$  loopers, leaving many hits in the CMS tracker. Vertex reconstruction was also attacked, and the code was optimized with autovectorization techniques, and by improving cluster splitting techniques.

#### 6.4.2 Technical improvements

Other changes introduced during the task force lifetime were technical, without any explicit changes in the algorithms’ code. We can cite

- transition to a newer compiler (from GCC 434 to 462),
- support for C++ extensions, like C++11,
- support for GCC autovectorization (which allows the use of vector instruction extensions in modern CPUs),
- libraries for memory allocation allowing for less RSS usage (after many tests, we chose JEMalloc),
- transition to a newer ROOT version, with less memory required in the I/O system (from 5.27 to 5.32)

#### 6.5 Performance of 2012 reconstruction code

The starting point with respect to performance is the release CMSSW\_4.4.X, used in the last period of 2011 data taking. When used to reconstruct data from the special High PileUp run, memory consumption (measured as RSS) increases past 1.5 GB before the event 100 is reconstructed, taking on average 80 seconds per event. The release which is already in production for 2012 data taking, CMSSW\_5.2.X, shows on the same events a memory consumption of 400 MB less, and takes 30 seconds per event. Memory use has hence decreased by 20-30%, with an outstanding factor 2.5 in events per second processed. With these figures, CMSSW reconstruction software is considered able to withstand 2012 LHC data taking conditions, even with an input rate from HLT exceeding the 300 Hz design requirements.

## 6.6 Alignment and Calibration

The high level of complexity and the large number of detector channels in CMS reflect in an elaborated structure for the management and computation of the detector calibration and alignment.

Most of the alignment and calibration workflows are fed with dedicated data samples, called AlCaReco, optimized both in terms of event selection and event content. Depending on the needs of the specific workflow, these samples can be selected offline, while performing the reconstruction, or directly online, at the High Level Trigger (HLT) level. The great flexibility of the HLT, which runs offline-quality software on a farm of commercial processors, is a key asset for this online selection since it guarantees an adequate rate of events that would not be selected by the standard trigger paths meant for physics analysis. An example of an online calibration stream is the one selecting events containing  $\pi^0$  and  $\eta$  candidates detected in ECAL and used for the inter-calibration of the PbWO<sub>4</sub> scintillating crystals. The calibration performance depends on the number of selected  $\pi^0$  candidates per crystal and on the signal to background ratio. The candidate di-photon decays are selected at the HLT level from events passing single- $e/\gamma$  and single-jet Level-1 triggers. After selection, only information about a limited region of ECAL (energy deposits in 20 to 40 individual crystals) near the  $\pi^0$  candidates is stored for the actual calibration. This allows to sustain a high rate of calibration events (1 to 10 kHz) whilst saving bandwidth and CPU time.

### 6.6.1 Prompt calibration loop

Conditions changing on a short time scale require a special calibration workflow designed to allow updates with very short latency. The handling of the data streams in the first step of the offline processing on the Tier-0 farm at CERN reflects this need and is organized as follows:

- express processing: reconstruction of a limited selection of data in order to give prompt feedback about the detector status and physics performance and to provide data for calibration workflows. The results of the express reconstruction for a given run are usually available one or two hours after the raw data are collected;
- bulk processing: reconstruction of the main data stream for physics analysis. This reconstruction step, also called prompt reconstruction, is delayed by 48 hours to allow for the computation of the fast-changing conditions. The output is divided in several Primary Datasets (PD) on the basis of the HLT paths used to select the events;
- calibration streams: streams of events selected at the HLT level and processed at Tier-0 for calibration purposes.

During normal operation of the CMS experiment about 300-400 Hz of data are processed in the bulk processing. Only a limited bandwidth, corresponding to about 10% of the bulk, is allocated for express processing in order to guarantee a fast reconstruction. A selection of data from the express and calibration streams is used to compute the updated conditions for a given run while the bulk of the data is buffered on disk. The calibration workflows run on a dedicated farm at CERN called the CMS Analysis Facility (CAF). In this way the prompt reconstruction can profit from the updated constants, reducing the need for offline reprocessing of the data. This workflow is called the prompt calibration loop (PCL) and is illustrated schematically in Fig. 1. The conditions currently updated through this kind of workflow are:

- measurement of the beam-line parameters;
- monitoring and masking of problematic channels of the silicon strip tracker to re-

spond to HV trips or noise;

- transparency corrections of the PbWO<sub>4</sub> crystals of the ECAL calorimeter.

Moreover, the delayed prompt reconstruction is also exploited to monitor possible movements of large structures of the silicon tracker, mainly due to thermal stress, and problematic channels in the electromagnetic and hadronic calorimeters allowing for quick reaction time in case of hot regions identified in the express reconstruction. The measurement of the three-dimensional profile of the luminous region where the LHC beams collide at CMS is an important component of the event reconstruction being used as an estimate of the primary interaction point prior to the reconstruction of the primary vertex. The position of the center of the luminous region and its width are determined using two independent methods with complementary systematic uncertainties.

The first uses the distribution of the reconstructed primary vertices to map the shape of the beam line; the mean three-dimensional position is determined with a 3D likelihood fit. The second method exploits the correlation between the transverse impact parameter and the azimuthal angle of tracks when the beam line is displaced from the expected position. With a sample of 1000 tracks, the position can be determined with a statistical precision of about 5  $\mu$ m. The fit is performed once per luminosity section (corresponding to 23 seconds of data taking) using tracks selected in the express stream. In a second step, ranges with stable parameters are collapsed, increasing the statistical precision and reducing the database storage size. Finally, the calibration object is validated and uploaded to the database. This allows the best possible knowledge of the position of the luminous region within a few hours of data being collected, with a time granularity that allows possible movements during the fill to be followed, as shown in Fig. 2.

## 6.7 Physics Validation

A fundamental aspect of the operations of a modern high energy physics (HEP) experiment is represented by the validation of the physics performance of its software. Several types of changes are regularly introduced in the chain that leads to the delivery of experiments software deployed at the computing centres for data processing and analysis. They can consist for example in algorithmic improvements in the objects reconstruction, simulation or trigger, in code performance optimisations or upgrade of basic components involved in the software building cycle like the compiler. The effect of these variations on the characteristics of the physics objects used for data analysis must be continuously and carefully assessed.

At least two products are necessary to achieve this goal: a high quality data quality monitoring (DQM) infrastructure allowing to collect in a collection of histograms the relevant properties of every dataset and a tool (RelMon) to compare different sets of histograms and estimate their level of compatibility.

During the years, the CMS collaboration developed a high quality DQM infrastructure for the certification of the acquired and Monte Carlo generated data [4]. The underlying principle is to append monitoring modules to the modules chain used for regular data processing. These modules allow to collect in histograms all the relevant quantities for dataset validation and data quality certification. The CMS Framework then takes care of writing such histograms in dedicated ROOT files. In order to efficiently make available to the physicists the enormous amount of histograms deriving from the operations of the CMS detector, DQM GUI servers are provided [5]. The ROOT files obtained running monitoring modules are automatically uploaded to the DQM servers which further process their content and store it in an internal database for performance reasons. A web interface can be accessed in order to browse all the



histograms contained in a given DQM server: the DQM GUI (figure 2).

Relmon is a general tool that allows to compare large sets of histogram pairs according to a statistical test defined by the user. The histograms that are analysed can be provided as ROOT histograms stored in ROOT files, also organised in a directory structure and the matching of corresponding histograms is done by name. The tests provided by RelMon to check the compatibility of two histograms can be: chi-squared, Kolmogorov-Smirnov and bin-to-bin. The first two are well known statistical procedures [3] and the tool relies on the their implementation provided by ROOT. The bin-to-bin test is not a traditional statistical test, but is useful in case the identity of two sets of histograms is to be checked. The aforementioned tests are considered to fail if their p-values are greater than a value set by the user. All the information about the amount of succeeding and failing comparisons is aggregated in a hierarchical way, closely following CMSSW's subsystems.

## 7 Computing

I. Fisk

## 8 Summary

T. Virdee

## References

- [1] R. V. e. A. Breskin, "The CERN Large Hadron Collider: Accelerator and Experiments, Vol. 2", *JINST* **3** (1998).
- [2] CMS Collaboration, "Internal Note CERN EN-EL",. EDMS 904391.
- [3] J. L. et. al., "The CMS Conceptual Design Report Part 2: Superconducting Coil Preliminary Design Report", *CEA Saclay report* **5C2100H-00001DA** (Sept. 1996) 6–14.
- [4] A. H. et. al., "The CMS Magnet Project Technical Design Report", *CERN/LHCC* **97-10** (May. 1997) 79–85.
- [5] r. t. C. c. T. Christiansen, "The CMS magnet test and cosmic challenge", in *IEEE 2006 Nuclear Science Symposium*. IEEE, 2006. arXiv:0805.1882.
- [6] CMS Collaboration, "CMS Physics Technical Design Report, volume I: Detector performance and software", *CERN/LHCC* **2006-001** (2006).
- [7] CMS Collaboration, "CMS technical design report, volume II: Physics performance", *J. Phys. G* **34** (2007) 995–1579, doi:10.1088/0954-3899/34/6/S01.
- [8] CMS Collaboration, "Particle-Flow Event Reconstruction in CMS and Performance for Jets, Taus and  $E_T^{miss}$ ", *CMS PAS* **PFT-09-001** (2009).
- [9] CMS Collaboration, "The CMS experiment at the CERN LHC", *JINST* **3** (2008) S08004, doi:10.1088/1748-0221/3/08/S08004.
- [10] V. Blobel, "Software alignment for tracking detectors", *Nucl.Instrum.Meth.* **A566** (2006) 5–13, doi:10.1016/j.nima.2006.05.157.

- [11] “Millepede II web page”, [https://www.wiki.terascale.de/index.php/Millepede\\_II](https://www.wiki.terascale.de/index.php/Millepede_II).
- [12] C. Kleinwort, “General Broken Lines as advanced track fitting method”, *Nucl.Instrum.Meth.* **A673** (2012) 107–110, doi:10.1016/j.nima.2012.01.024, arXiv:1201.4320.
- [13] V. Blobel, C. Kleinwort, and F. Meier, “Fast alignment of a complex tracking detector using advanced track models”, *Comput.Phys.Commun.* **182** (2011) 1760–1763, doi:10.1016/j.cpc.2011.03.017, arXiv:1103.3909.
- [14] C. Kleinwort and F. Meier, “Alignment of the CMS Silicon Tracker - and how to improve detectors in the future”, *Nucl.Instrum.Meth.* **A650** (2011) 240–244, doi:10.1016/j.nima.2010.11.187, arXiv:1010.2039.
- [15] D. Brown et al., “Local Alignment of the BABAR Silicon Vertex Tracking Detector”, *Nucl.Instrum.Meth.* **A603** (2009) 467–484, doi:10.1016/j.nima.2009.02.001, arXiv:0809.3823.
- [16] M. Swartz et al., “A new technique for the reconstruction, validation, and simulation of hits in the CMS pixel detector”, *PoS VERTEX2007* (2007) 035.
- [17] M. Swartz, “CMS pixel simulations”, *Nucl.Instrum.Meth.* **A511** (2003) 88–91, doi:10.1016/S0168-9002(03)01757-1.
- [18] R. Frühwirth, “Application of Kalman filtering to track and vertex fitting”, *Nucl. Instrum. Meth. A* **262** (1987) 444, doi:10.1016/0168-9002(87)90887-4.
- [19] W. Adam et al., “Reconstruction of electrons with the Gaussian-sum filter in the CMS tracker at LHC”, *ECONF C0303241* (2003) TULT009, doi:10.1088/0954-3899/31/9/N01, arXiv:physics/0306087. [J.Phys.G31:N9,2005].
- [20] T. Speer et al., “Vertex Fitting in the CMS Tracker”, *CMS NOTE* **2006/032** (2006).
- [21] K. Rose, “Deterministic Annealing for Clustering, Compression, Classification, Regression and related Optimisation Problems”, *Proceedings of the IEEE* **Vol. 86, Issue 11** (1998).
- [22] T. Miao et al., “Beam Position Determination using Tracks”, *CMS NOTE* **2007/021** (2007).
- [23] CMS Collaboration, “CMS Physics Technical Design Report”, *CERN/LHCC* 2006-001 (2006).
- [24] CMS Collaboration, “Performance of the CMS drift-tube chamber local trigger with cosmic rays”, *J. Instrum.* **5** (2010) T03003, doi:10.1088/1748-0221/5/03/T03003, arXiv:0911.4893.
- [25] P. Arce et al., “Bunched beam test of the CMS drift tubes local muon trigger”, *Nucl. Instrum. Meth. A* **534** (2004) 441.
- [26] CMS Collaboration, “The Level-1 Trigger, Technical Design Report”, *CERN/LHCC* 2000-038 (2000).
- [27] E. Gatti et al., “Optimum geometry for strip cathodes or grids in MWPC for avalanche localization along the anode wires”, *Nucl. Instrum. Meth.* **163** (1979) 83–92.

- [28] J. Gordon and E. Mathieson, "Cathode charge distributions in multiwire chambers: I. Measurement and theory", *Nucl. Instrum. Meth.* **227** (1984) 267–276.
- [29] J. Gordon and E. Mathieson, "Cathode charge distributions in multiwire chambers: II. Approximate and empirical formulae", *Nucl. Instrum. Meth.* **227** (1984) 277–282.
- [30] CMS Collaboration, "The Muon Project Technical Design Report", *CERN/LHCC* 97-32 (1997).
- [31] CMS Collaboration, "Performance of the CMS Drift Tube Chambers with Cosmic Rays", *J. Instrum.* **5** T03015 (2010).
- [32] CMS Collaboration, "Performance of CMS muon reconstruction in cosmic-ray events", *JINST* **5** (2010) T03022, doi:10.1088/1748-0221/5/03/T03022, arXiv:0911.4994.
- [33] CMS Collaboration, "Measurement of tracking efficiency", *CMS Physics Analysis Summary* **CMS-PAS-TRK-10-002** (2010).
- [34] CMS Collaboration, "Particle-flow event reconstruction in CMS and performance for jets, taus, and missing  $E_T$ ", *CMS Physics Analysis Summary* **CMS-PAS-PFT-09-001** (2009).
- [35] CMS Collaboration, "Commissioning of the particle-flow event reconstruction with leptons from J/Psi and W decays at 7 TeV", *CMS Physics Analysis Summary* **CMS-PAS-PFT-10-003** (2010).
- [36] CMS Collaboration, "Prompt and non-prompt J/ $\psi$  production in pp collisions at  $\sqrt{s} = 7$  TeV", *Eur. Phys. J. C* **71** (2011) 1575, doi:10.1140/epjc/s10052-011-1575-8.
- [37] CMS Collaboration, "Measurements of inclusive W and Z cross sections in pp collisions at  $\sqrt{s} = 7$  TeV", *J. High Energy Phys.* **01** (2011) 080, doi:10.1007/JHEP01(2011)080.
- [38] CMS Collaboration, "Measurement of the inclusive W and Z production cross sections in pp collisions at  $\sqrt{s} = 7$  TeV with the CMS experiment", *J. High Energy Phys.* **10** (2011) 132, doi:10.1007/JHEP10(2011)132, arXiv:1107.4789.